Title: Evidence of Completion of Milestone 3: Experimental Testbed

Author(s): Geller, Drew Adam

Intended for: Report

Issued: 2021-06-29

ARPA-E Open 2018: Award Number 1858-3663
PI: Johanna Mathieu, University of Michigan

Evidence of Completion of Milestone 3: Experimental Testbed

Construction and deployment of the experimental testbed encountered scheduling delays, and completion was subsequently further delayed by curtailment of LANL operations due to the COVID-19 pandemic. Between March and October, LANL was in a state of "Limited Operations" such that only certain "mission-critical" work was performed with approximately 25% on-site staffing. In spite of this, much of the lab space preparation was executed and the five initial model house" units[1] were constructed and installed during this period. Still, the pace of construction was impacted by difficulty in coordinating personnel and minimizing contact between workers.

Currently, LANL is in the mode of "Normal Operations with Maximum Telework," which continues to limit the availability of on-site personnel. Nevertheless, Milestone 3 is now complete. This document describes the experiment as it is deployed and provides status for each task for this milestone.

Milestone 3 is given in the SOPO as being due in quarter 4 (ending 6/10/2020) and is described thus:

> *Milestone 3: Experimental testbed validated with Pecan Street data (LANL, Geller)*
>
> *Experimental testbed with 20 physical TCLs and a fully configurable (in terms of connectivity/quality) communication network used by the load controller. This testbed will be coupled to the simulation testbed, which will simulate additional loads and the distribution network.*
>
> *We will report our detailed experimental plans and photos of experimental system.*
>
> *The experimental testbed should accurately capture TCL real and reactive power consumption during normal operation to within 5% RMSE error with respect to actual TCL data collected by Pecan Street.*

The validation of this milestone will consist of (1) credible evidence of the experimental plans and photos of the experimental system and (2) an analysis to show whether the experimental testbed accurately captures TCL real and reactive power consumption during normal operation

[1] In this report, "model house" refers to an insulated, modular assembly that contains a heat source and an air conditioner, the latter being the thermostatically controlled load (TCL) operating in the experiment. The control algorithms will control all the air conditioners and use measurements of their power consumptions and the temperature in their model houses to benchmark performance of the control.

to within 5% RMSE error with respect to actual TCL data collected by Pecan Street. Because this milestone is delayed, the following sections describe the progress on each component task.

This table summarizes progress towards the milestone for each subtask and highlights the inter-dependencies among them:

M3: Experimental testbed validated with Pecan St data	**Due Date:** 6/11/2020	**% complete:** 100% (from 90% last quarter)
Comprises Experiments, tasks 2.1-2.7		
2-1 Prepare laboratory space		100%
Laboratory space cleared and prepared to provide electrical power and structural support as needed. *Specifically,* detailed facility requests drafted to initiate the engineering work. Scoping meetings with engineering and crafts held to determine the design. Hazards assessed, work authorization documents developed, and modifications scheduled and executed.		
Performer Input: We have cleared the laboratory space and provided electrical power and structural support as needed to field the 20-TCL experimental testbed. The lab is operational.		
2-2 Design and fabricate 5 trial TCLs		100%
Five trial TCLs (air conditioners) designed and constructed. *Specifically,* bill of materials for mechanical components and supports, electrical components, controls, sensors, and cabling created. Sketches drafted for construction. Components for 5 TCLs purchased, and TCLs assembled and installed in the lab. Instruments and sensors calibrated as necessary.		
Performer Input: We created bills of materials for mechanical components, electrical components, controls, sensors, and cabling. Sketches were drafted for construction. Components for 5 TCLs were purchased. Instruments and sensors for the 5 TCLs were calibrated as necessary. We assembled and installed 5 TCL units in the lab.		
2-3 Design and fabricate communication network for control signals, state/output feedback signals		100%
Aggregator-to-TCL communication network for control signals and state/output feedback designed and programmed. *Specifically,* data acquisition system and data flows for preliminary testing and logging developed. Computer, software, and I/O interfaces purchased. Data acquisition system assembled, tested, and validated.		
Performer Input: We have developed a data acquisition system and mapped the data flows for the experiment. We have procured a computer, software environment, and I/O interfaces. The data acquisition system is operational, has been tested with the 5 trial TCLs, and has been validated through experiments and testing hardware.		
2-4 Conduct preliminary testing to validate testbed against Pecan St data		100%
Trial TCLs tested to validate the testbed against PSI data. *Specifically,* experiments performed spanning a range of mean temperatures and heat loads. Phase locking and the impact of TCL heterogeneity investigated. (Scaled) power data compared against PSI data for validation.		
Performer Input: The 5 trial TCLs were tested for a range of temperatures and programmable heat loads. Temperature deadband was varied. Phase locking was evaluated for free-running TCLs. Tools were developed to examine TCL heterogeneity, and data were compared against PSI field test data.		
2-5 Review/revise design based on validation results		100%
Deviations of the trial TCLs from the PSI data investigated. Insulation, positioning, heat load range, and/or ventilation changed as necessary to minimize deviations from the PSI data.		

Performer Input: The experiment's distributions of cycle durations, duty cycles, and power transients were compared against PSI data. Modifications were made to increase the range of duty cycles possible and to ensure cycle durations in agreement with PSI data.	
2-6 Fabricate remaining 15 TCLs	100%
Remaining 15 TCLs fabricated. *Specifically,* final sketches from trial TCLs reviewed/revised and supplies for the remaining TCLs purchased. TCLs fabricated, transported, and installed in the lab space. Data acquisition system expanded to capture the extra channels. Instruments and sensors calibrated as necessary. Hooks for the simulation testbed to interact with experimental testbed and data acquisition system designed and implemented.	
Performer Input: All signal hardware to support additional TCLs is complete. Hooks for simulation testbed to interact with experimental testbed are being developed. Supplies were purchased for building the remaining units. All 20 units are complete and installed. Interchangeable sensors were selected, where possible, to eliminate calibration time. Hooks were provided for the simulation testbed to interact with the experimental testbed.	
2-7 Conduct comprehensive testing to validate testbed against Pecan St data	100%
The full experimental testbed comprehensively tested and validated against the PSI data. Specifically, data from experiments on the complete testbed compared with the data from the trial TCLs. Data examined for phase locking or coupling. Performance assessment completed, specifically, whether (scaled) TCL real and reactive power consumption during normal operation in the experimental testbed is within 5% RMSE error with respect to the PSI data.	
Performer Input: Experiments to benchmark the complete testbed are finished. Experiments were run at different temperature set points, internal heating rates, and temperature deadbands.	

The following subsections enumerate the tasks required to complete Milestone 3 and provide the credible evidence for completion in each task. Milestone 3 encompasses the completion of the Tasks 2.1-2.7 as described below.

2.1 Laboratory Space Prepared

The SOPO describes Task 2.1 (Laboratory space prepared) as follows:

> *Laboratory space cleared and prepared to provide electrical power and structural support as needed.*
>
> *Specifically, detailed facility requests drafted to initiate the engineering work. Scoping meetings with engineering and crafts held to determine the design. Hazards assessed, work authorization documents developed, and modifications scheduled and executed.*

This task is complete. Protection of the transformer and electrical panel was provided due to the co-located operation of motorized forklifts and vehicles. This delayed closeout, but the system is now energized.

Although the warehouse is quite large, the space is shared by three organizations that store large industrial items and hardware. For this reason, a compact configuration of the model houses was necessary, if not ideal with respect to thermal isolation between TCL units. The

model houses are stacked in pairs, using pallet racks arranged in a U-shape. Due to seismic requirements, it was necessary to anchor the racks. This, in turn, required a ground penetrating radar (GPR) survey before installation. Due to the nature of the flooring surfaces, 2”x8” wood beams were used under the footings to spread the weight of the assemblies. The complete testbed occupies a space approximately 250 sq. ft. in area plus another 200 sq. ft. to accommodate the electrical transformer, panelboard, and margins.

Each of the 20 model houses requires its own 20 A, 120V circuit, and the density of these circuits is much higher than that found in commercial buildings, which is typically around 1 VA/sq. ft. [1] To do this, a 480 V line of sufficient amperage was located for transforming to 120 V in three phases. Unfortunately, the electricians found that there was a missing line to the upstream circuit breaker, and so they had to install new wiring to feed the 480V:120V transformer.

Because it was not originally known what power meter would be used in the measurements, the installation of an eGauge at the panelboard was not included in the original plan. A field change was implemented to add the eGauge in an enclosure, with a three-phase circuit breaker feeding the power and phase inputs. The eGauge enclosure is near, but not within, the pallet racks so extensions of 18 AWG twisted pair leads will be used to connect the current transformer (CT) sensors to the line splitter at each TCL, following eGauge’s recommendation [2].

The evidence for completion of this task is:

- Image of the workspace with installed support racks and electrical facilities
- Image from drawing set for the engineering change package

Figure 1 is a selection from the engineering drawings depicting the arrangement of rack supports as seen from above, and Fig. 2 shows the assembly before installation of the model house units with TCLs.

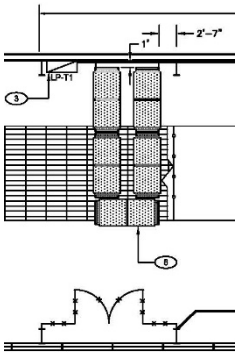

Figure 1: The model houses are arranged on two levels of shelving, with the shelving arranged in a U-shape for compactness. The section of the warehouse allotted to the experiment is approximately 20’ by 25’. The new panelboard supporting the convenience outlets is located in the upper left of the image. The transformer, not shown, is located to the left of the panelboard.

Figure 2: The completed structural and electrical assembly prior to installation of testbed modules. The eGauge power meter is mounted inside a small NEMA enclosure attached to the vertical beam on the left hand side of the photo.

2.2 Five Trial TCLs Designed and Fabricated

The SOPO describes Task 2.2 (5 trial TCLs designed and fabricated) as follows:

> *Five trial TCLs (air conditioners) designed and constructed.*
>
> *Specifically, bill of materials for mechanical components and supports, electrical components, controls, sensors, and cabling created. Sketches drafted for construction. Components for 5 TCLs purchased, and TCLs assembled and installed in the lab. Instruments and sensors calibrated as necessary.*

This task is complete. The 5 trial model houses with TCLs were fabricated and installed in the laboratory. The units are instrumented for temperature and electrical power measurements.

Evidence of completion is provided in:

- Mechanical bill of materials
- Electrical bill of materials
- Schematic of hydronic loop
- Picture of voltage to current converter for programming heater
- Schematic of heater and compressor control box
- Example data from thermistor
- Photographs of the completed assembly
- Calibration data for the programmable heat load

Bills of materials for mechanical and electrical parts, containing most important components, are given in Appendix B.

The controlled load in these model houses is the compressor of a window air conditioner (AC). Because we are interested in developing controls and not in coping with or analyzing software built into modern "smart" air conditioners, an important selection criterion was that the AC have a simple mechanical thermostat. These thermostats act as simple switches to turn the AC's compressor on/off and can readily be replaced with a relay that can be computer controlled such that the software now provides the thermostat function. Another important restriction on our choice was the refrigerant used in the unit. LANL is currently capable of recovering and disposing of R410A but does not accept the R32 refrigerant used in newer ACs. R32 has a lower impact as a greenhouse gas and does not cause ozone depletion, but it is slightly flammable. Considering these criteria, we selected the LG model LW5016 as a test system (see the cut sheet in Appendix A).

A heat source is needed for emulating heat leaks as from solar irradiance or from internal heat sources such as lighting. Another heat source to consider is that of human occupancy in conditioned spaces. These loads are not part of the TCL in terms of being controlled by the control algorithms, and so they are not measured by the power meter. But they are incorporated in the model house units, and the powers are calculated, programmed, and logged by the experiment control software itself. Some randomization of this aggregate heat load is useful for breaking the periodicity of the on/off cycles in the relatively constant (and unrealistic) laboratory environment. The programmable power source is provided using a silicon controlled rectifier (SCR) with variable time base, zero-cross switching. A chassis is located at each TCL to provide this modulated heater power and also to house a TTL- or CMOS-programmable relay circuit to turn on/off the air conditioner's compressor under computer control. A sketch of this chassis is shown in Fig. 3, and a picture of a simple relay board ready for installation is given in Fig. 4.

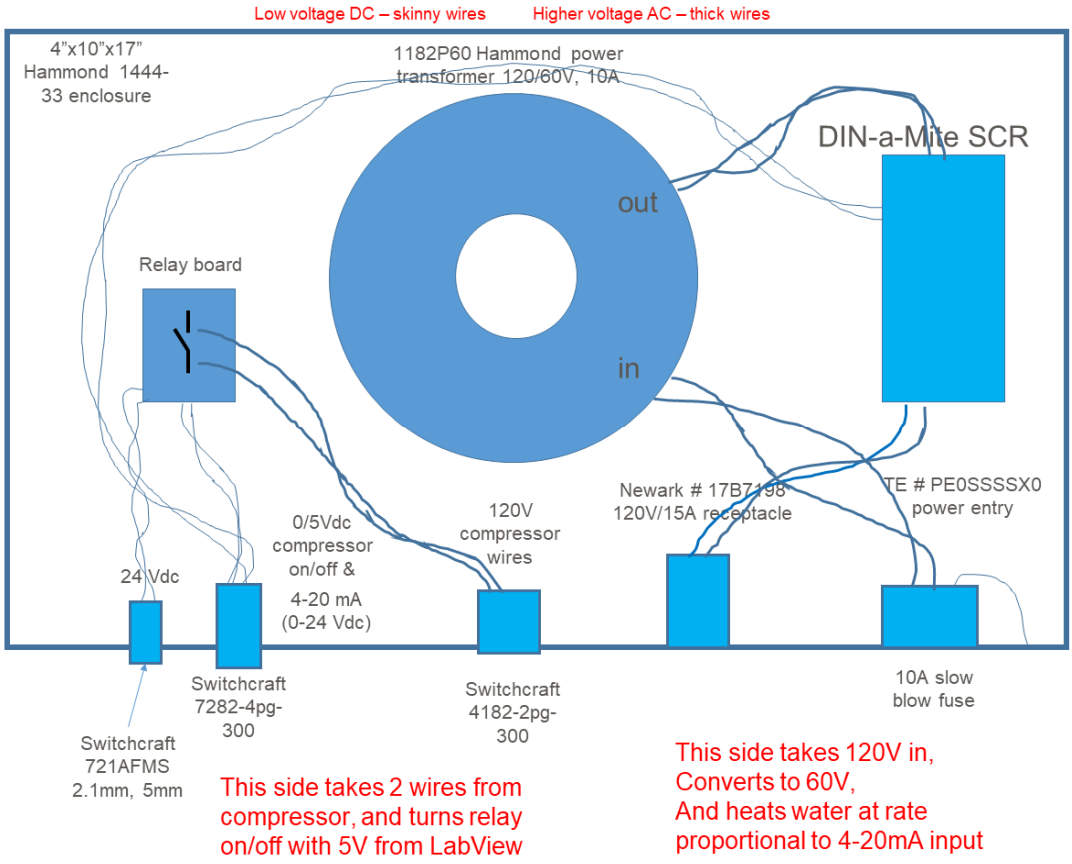

Figure 3: A chassis located at each TCL manages power to the water heater and the on/off state of the air conditioner's compressor, driven by 4-20 mA signals and 5 V logic signals from the data acquisition system, respectively.

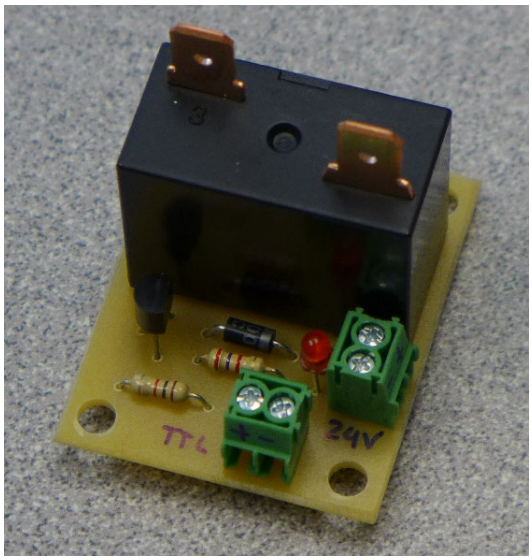

Figure 4: The motor-control relay module for switching the AC's compressor.

The output heater power is programmed with a 4-20 mA signal through two terminals of the SCR device. The highest channel density devices available for the National Instruments data acquisition system have voltage outputs, so we use voltage programming instead of current outputs. The 4-20 mA loop currents are produced from 1-5 Vdc analog output signals using a converter circuit based on the Texas Instruments XTR110. A 10-channel interface board is shown in Fig. 5.

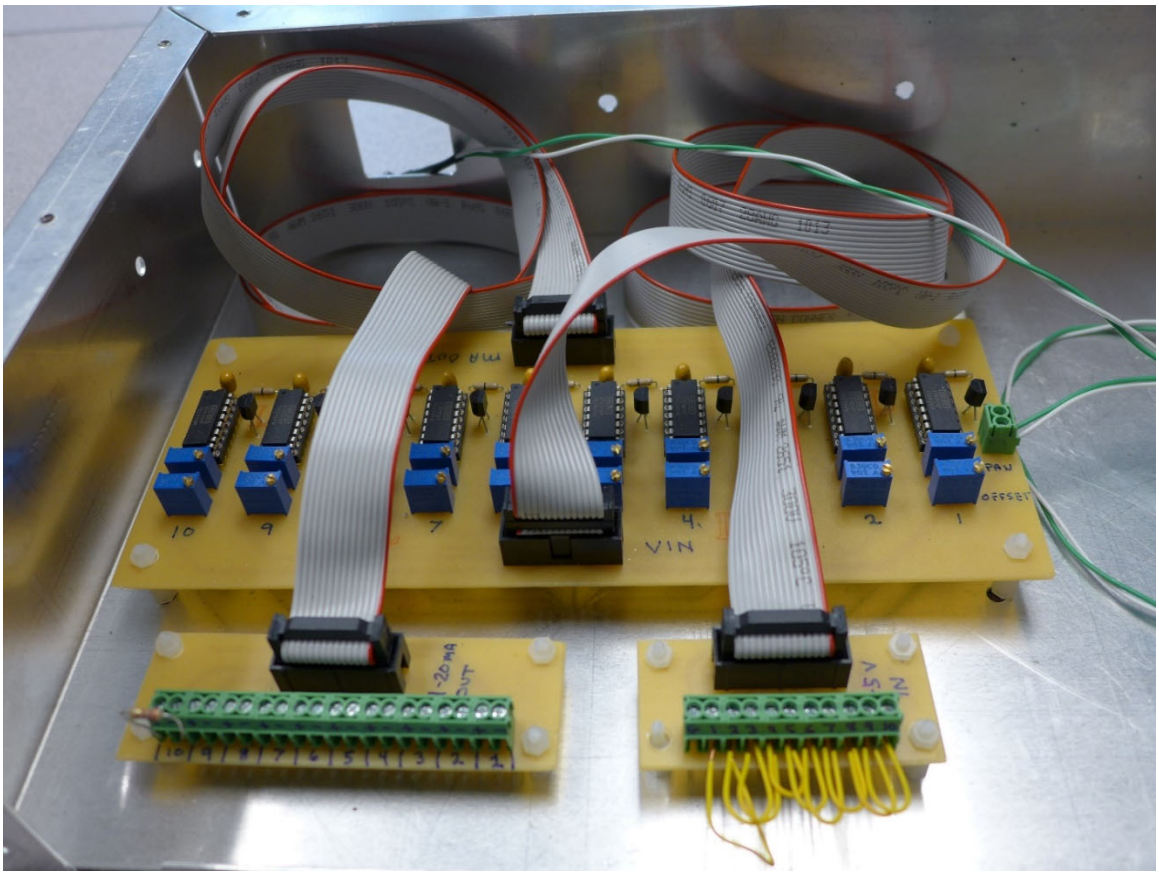

Figure 5: Voltage-to-current converter board with 10 channels. The inputs on the right daughter board are shown jumpered together so that all channels can be tested with the same 1-5 V analog input signal. Isolation between the individual input channels is also checked before use.

For a realistic system with an AC-based TCL to control, similar to a house or office space, there must be significant heat capacity emulating the furniture, equipment, drywall, and occupants. So rather than directly heating the air, we inject heat into a large volume of water, 20 or 30 gallons, in a point-of-use electric water heater. The water heater is heavily insulated, and the heat capacity of the solids is difficult to determine. However, the incorporation of the water heater in a hydronic loop allows for adjustable thermal conductance between the water and air, at the expense of two additional constant-rate heat sources in the form of the water pump and of the fan circulating air through the heat exchanger. A sketch of the hydronic loop is provided in Fig. 6. Two sizes of water heater are used, 20 and 30 gallon, in order to have a diversity of

effective heat capacities and a wide range of natural cycle durations. The larger tanks make longer cycle durations possible, for a given set of thermal conditions, but it is difficult to adjust the hydronic loop's flow rate low enough for the shorter cycle durations of the model houses with 20 gallon tanks.

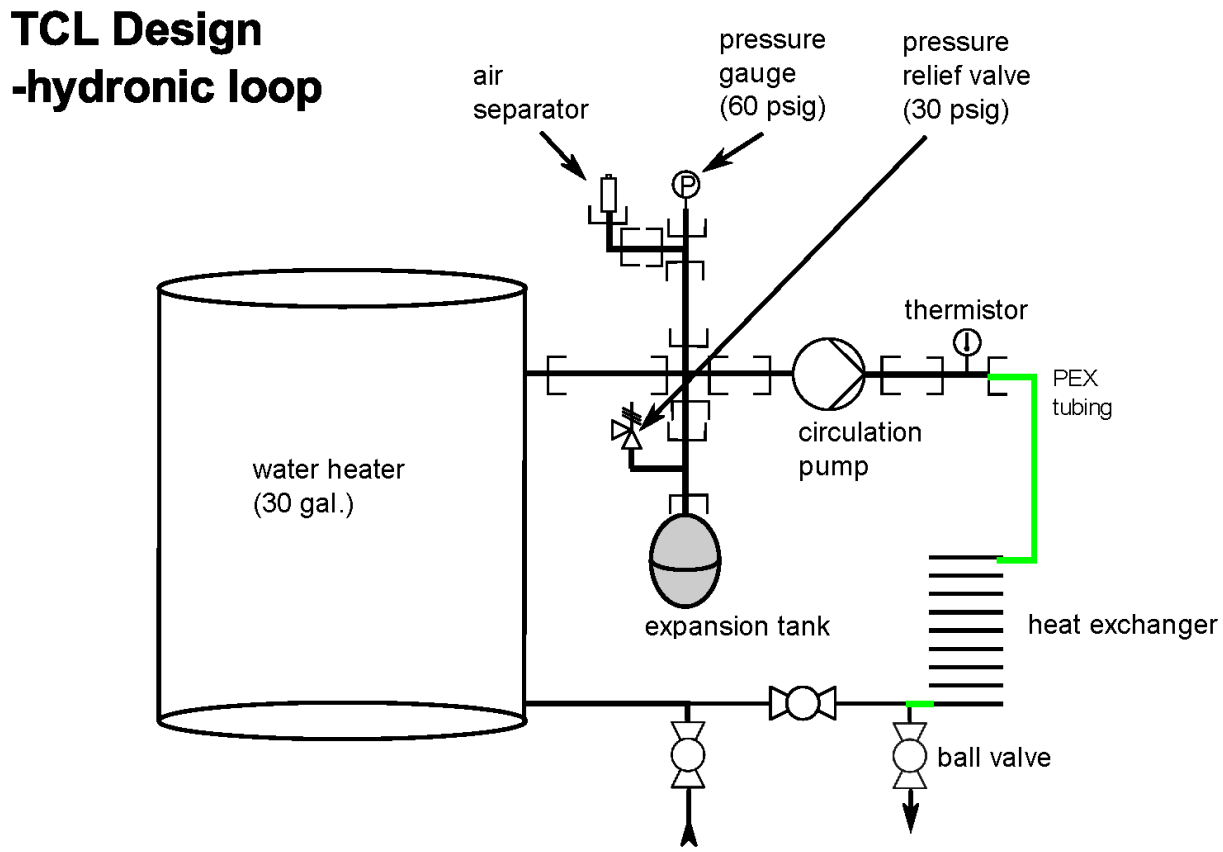

Figure 6: Hydronic loop for stochastic or user-programmed heating loads into the TCL.

The original concept included the use of thermocouple probes for measuring temperatures throughout the testbed, due to their rugged construction and the availability of standard National Instruments interface cards configured to read them directly. A set of K-type thermocouples was calibrated with a

Fluke 9102S Dry Well for this purpose. However, the voltage of a K-type thermocouple only changes by about 40 µV/°C [3] around room temperature. Combined with the likelihood of other electronic noise, the resolution of the temperature sensor would be around 1 °C without filtering and signal averaging over a 1 s window. Instead, we procured a set of 10 kΩ negative-temperature-coefficient (NTC) thermistors, some of which are sheathed in 1/8" NPT fittings for easily immersing the probe in the hydronic system to monitor water temperature. In contrast to thermocouples, thermistors are low cost, highly sensitive, and commonly used in ordinary thermostats. Although thermistors are highly non-linear in their resistance-to-temperature characteristic, the temperature is easily calculated within a formula node in LabView so that no complicated electronics are needed to linearize the signal. Also, thermistors can be highly reproducible in their temperature characteristic. All of the sensors procured so far are specified by the manufacturer to conform to their Type-J response curve [4] and to be interchangeable to an accuracy of 0.1 °C.

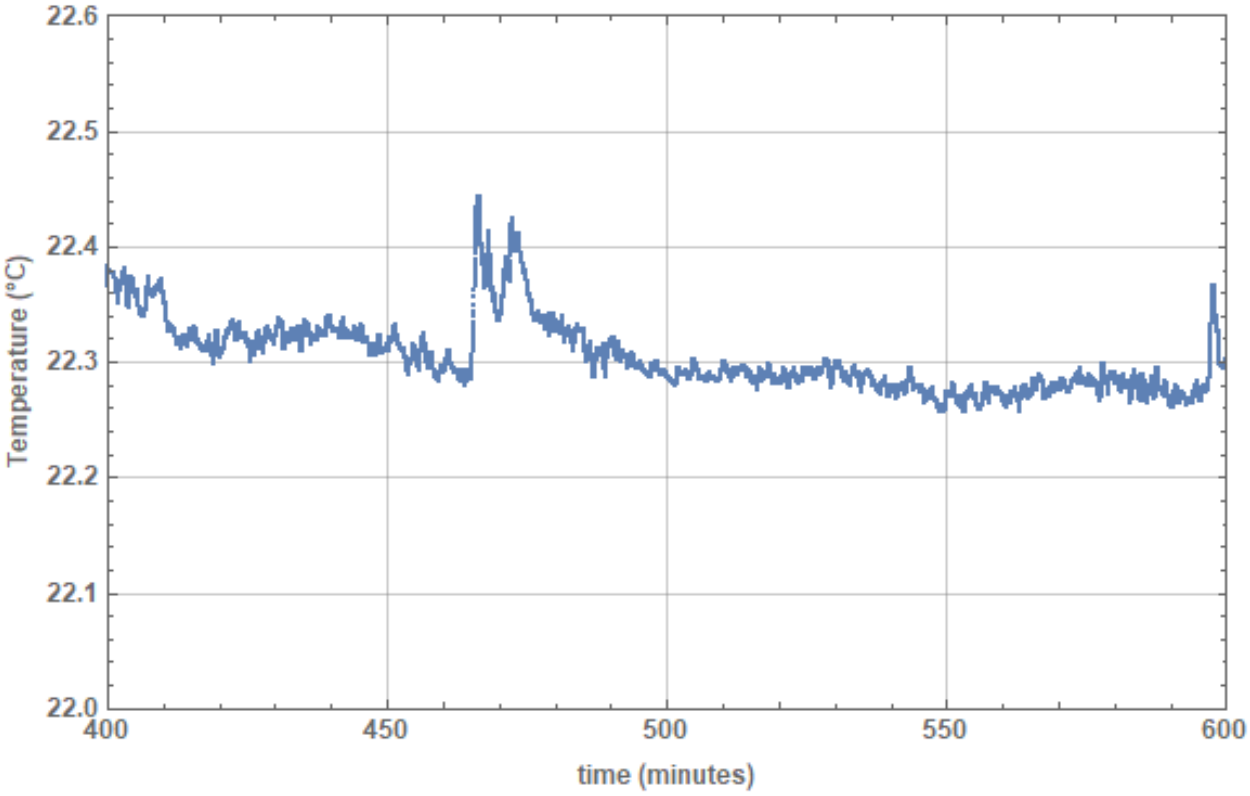

Figure 7: Testing of thermistor noise and resolution.

A brief test of one bare thermistor is shown in Fig. 7, where the resistance is measured simply as one arm of a voltage divider and the sensor was taped to a thin-wall metal tube. The large spikes occurred when the door next to the workbench was opened while the test was running. The signal has been averaged at 1000 Hz over 1 s, and the noise in temperature measurements is about 0.01 °C, which should be adequate resolution for the testbed.

One downside of thermistors compared to thermocouples is that thermistors may read incorrectly due to self-heating by the drive currents used to measure them, depending on the heating rate and thermal contact to the environment [5]. In the bench test here, it was found that the self-heating offset was only ~0.02 °C for the 25 µW of heat generated with test voltages of 0.5 V near room temperature.

Incorporating this hardware, the initial five trial TCLs have been built and installed, as is shown in Fig. 8.

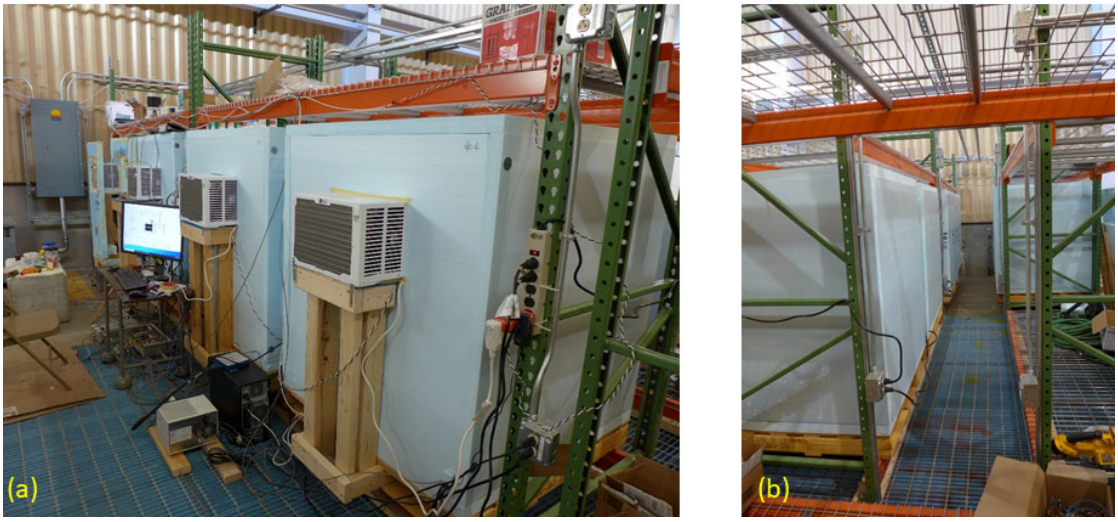

Figure 8: (a) shows four of the model houses with TCLs installed along one side of the structure. The fifth and sixth units are seen on the right side of (b) at the end of the next row of racks. The foam access panels are in place, completing each unit's thermal enclosure. Additional foam and tape will be used to improve thermal insulation from the lab space.

The heat load inside the model house comes primarily from

(1) Programmable heating of the water heater,
(2) Power into the water circulation pump,
(3) Power into the air circulation fan under the air-water heat exchanger, and
(4) Heat leaks through the foam walls.

Exact values are not needed, but it is useful to have baseline values for each unit as a check of consistency, for modeling, and to estimate duty cycles for each unit. These estimates can be used to program the desired duty cycles in experiments. The pump and fan loads were measured using the installed eGauge power meter with the CT sensor at each station:

Unit	**Fan (W)**	**Pump (W)**
1	136.9 ± 0.2	80.0 ± 0.2
2	43.9 ± 0.2	82.7 ± 0.2
3	44.5 ± 0.2	86.4 ± 0.2
4	44.7 ± 0.2	84.3 ± 0.2
5	44.0 ± 0.1	84.2 ± 0.2

Unit #1 differs from the others because it was the prototype unit and has a different air circulator. This unit was used in part because the AC in unit #4 failed and was replaced with the AC from unit #6, the last of the new TCLs. The unit #1 also has a different AC, which contributes to inhomogeneity of the systems.

The local power chassis (containing the circuit of Fig. 3) were also checked for linear output of the power to the water heater versus the power requested by the computer. This was checked for all units using the eGauge; and for unit #4 the power delivered was also checked against a second meter, a Hameg HM8115-2, which agreed with the eGauge. Data from this testing is shown in Fig. 9.

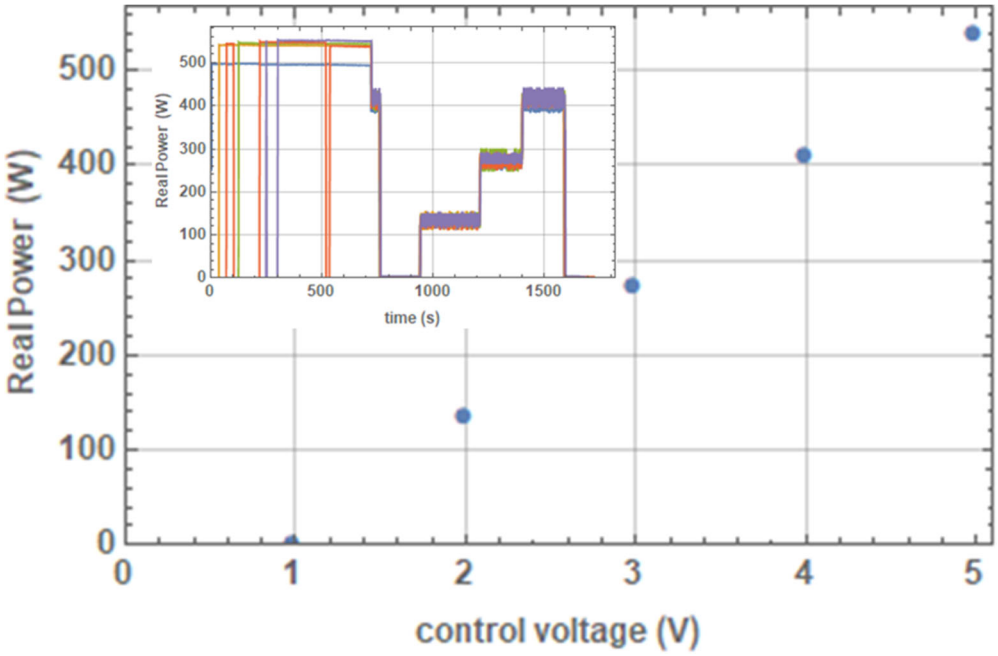

Fig. 9: Time-averaged power into the water heater vs. control voltage for unit #4 is linear as expected. The 1-5 VDC control voltage produces 0-520 W of heat in the model house. A mean heating rate and variable noise may be selected to emulate occupancy or to bias the duty cycle of the fixed-speed AC. The inset shows the stepwise power test data for all 5 water heater units. The noise on the middle steps is due to gating by the power control SCR and is about ± 12 W there, whereas the uncertainty is ± 0.4 W at the limits of the range where the AC voltage is fully passed or blocked. Gating patterns of the SCR do not necessarily fill the averaging time of the power meter with the same number of 120V cycles in each measurement interval, for intermediate power settings.

2.3 Communications Network Designed and Fabricated

The SOPO describes Task 2.3 (communications network designed and fabricated) as follows:

> *Aggregator-to-TCL communication network for control signals and state/output feedback designed and programmed.*

Specifically, data acquisition system and data flows for preliminary testing and logging developed. Computer, software, and I/O interfaces purchased. Data acquisition system assembled, tested, and validated.

This task is complete. Evidence of completion is provided in:

- Data acquisition components table
- Diagram of the data flows
- Data from real TCLs demonstrating software thermostatic control

A bill of materials describing the components of the data acquisition system is given in Appendix C. For fast data acquisition capability, National Instruments components were selected. The PXI express interface was chosen due to its high channel density and ease of expansion. In general, these devices can be programmed in C, Visual Basic, or MATLAB, but the fastest and most flexible development environment for varied devices is the LabView Developer package which is used here for the experiment control software.

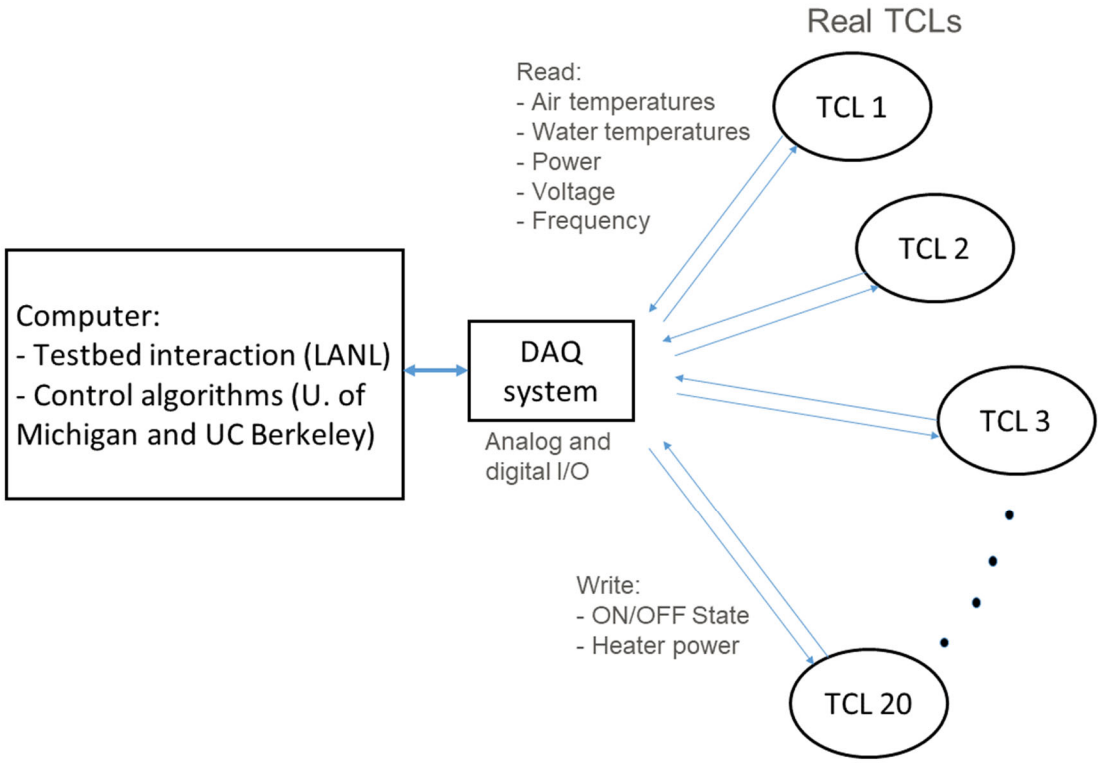

Figure 10: A single desktop computer controls the experimental testbed through a data acquisition system. All model houses are connected to LabView hardware for thermal and control signal I/O. Power, voltage, and frequency are acquired from an eGauge power meter via TCP/Modbus.

There are three basic modes in which the software must control the testbed for each experimental sequence. In the first mode, given a starting condition or set of temperature set points for the ensemble, the software must control the testbed for a long enough time with simple fixed-deadband thermostatic control for all TCLs to reach steady-state operation. Then, in the second mode, the testbed is managed with this simple thermostat control or with external control requests generated by the control software (developed by UM and UCB). In the former case, the experiment runs freely for a fixed period of time, and the data is used to tune heat exchange if necessary for the TCL to operate with cycle duration consistent with the Pecan Street field data. The fixed length of this mode is useful for allowing replay of schedules for the heat input, if desired, for comparison of different control mechanisms or thermostat settings. In the third and final mode, ordinary thermostatic control resumes for an unlimited duration so that the operator can safely shut down the system when convenient.

The data acquisition runs on a general-purpose computer rather than within a real-time environment. This is done for convenience in development and for fielding the UM/UCB control software (see Fig. 10). Therefore, the several-millisecond delays for operating system process executions and disk access could potentially cause software timing of the steps of the control loop to have time deviate and, worse, dilate from the intended data acquisition rate. This would cause problems with synchronizing the physical experiment to the simulation testbed and controls. Fortunately, NI's LabView has a function for access to the computer's millisecond counter, so that the average step timing of the main loop can be maintained as one cycle per second (limited by the minimum data acquisition rate of the eGauge power meter). That is, the execution of a given cycle may lag the 1000^{th} tick of the computer's hardware timer by x milliseconds, but the next iteration will again be armed on the 1000^{th} tick rather than being delayed by x ms. This serves to keep measurements on schedule, and the timestamp of each loop step is compared against the timestamp from the loop's start to ensure this.

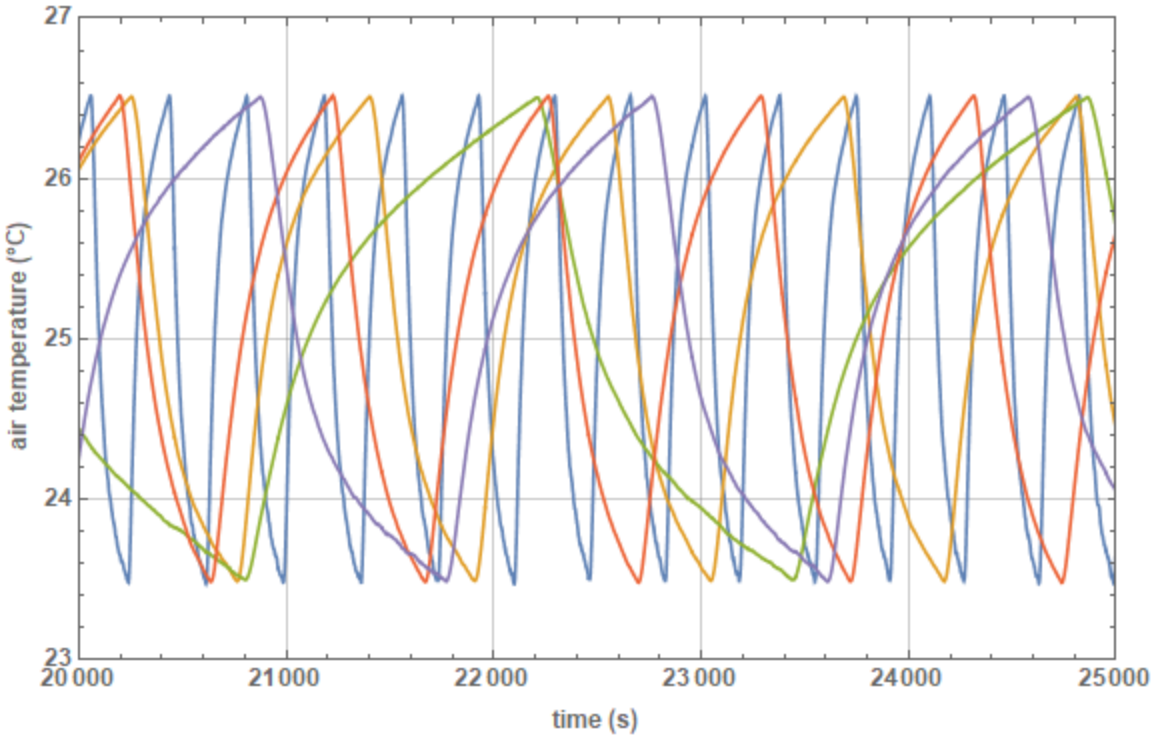

Figure 11: Air temperature data during a typical free-run experiment. Each color represents air temperature inside one of the five initial physical model houses of the experimental testbed. The computer system implements the thermostatic control, sets the local heating rate in each unit, and collects and logs all temperature, voltage, and power data. Requests to periodically switch the cooling ON or OFF in each unit can also be generated within the code, and these requests are fielded by a subroutine that determines whether the requests can be enacted based on the current temperatures and the times since the last state changes of each TCL.

The first version of functioning LabView data acquisition code is complete and has been used to collect preliminary data from the testbed (see Fig. 11). Bugfixes and new features are handled with updates to the code. For example, a feature was added to separately switch on/off power meter readings during the run, in order to disconnect or reboot the meter if required without communication timeouts interrupting testbed operation. Because testbed experiments may frequently need changes to the control settings recipe, the software reads and checks the input recipe and then—if no errors are detected—writes the settings to a timestamped log file so that they can be easily connected to the data file of the corresponding run.

Voltage connections to the thermistor sensors were made using Belden #8723 shielded signal cable, with care taken to ground the shield and avoid ground loops. The shields were grounded only to the National Instruments PXIe-4302 device, while the negative differential voltage was derived from the separate ground of the voltage divider's power supply at the voltage divider breakout box. All thermistors are driven from a single benchtop power supply. A selection of voltage points was checked with a portable meter to verify readings to and from the data acquisition devices.

2.4 Preliminary Testing and Validation Conducted

The SOPO describes Task 2.4 (preliminary testing and validation conducted) as follows:

> *Trial TCLs tested to validate the testbed against PSI data.*
>
> *Specifically, experiments performed spanning a range of mean temperatures and heat loads. Phase locking and the impact of TCL heterogeneity investigated. (Scaled) power data compared against PSI data for validation.*

This task is complete. The five trial model houses were tested to control a range of air temperatures (22-28 °C) inside the foam boxes. The temperature deadband for this control was varied between 2-4 °C, and variable heating rates of 0-520 W were injected to the houses through the water heater. The results of open-loop testing are reviewed in this section and compared against field test data from PSI.

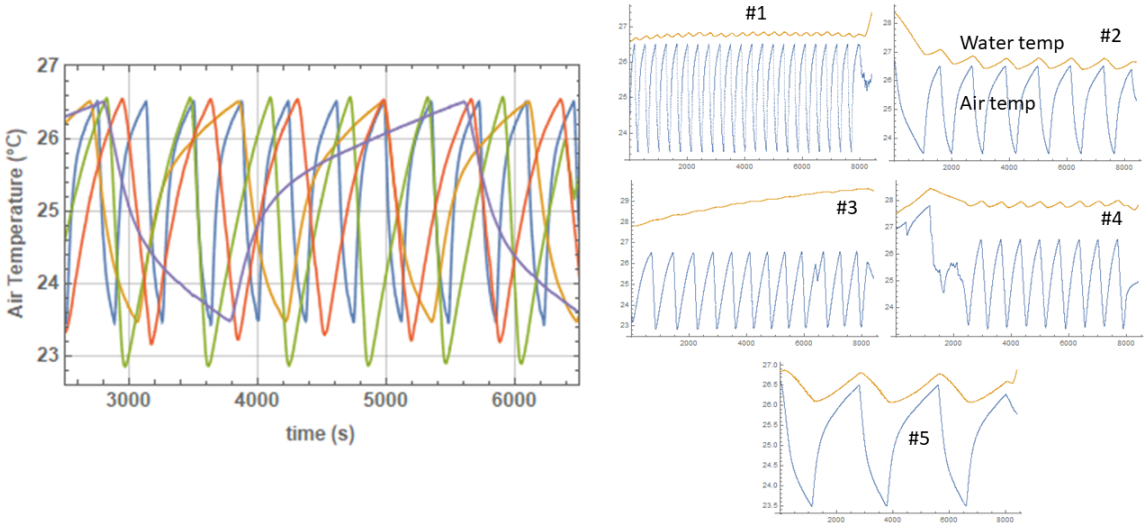

Fig. 12: On the left, the air temperatures of all five units are plotted together. The model houses run independently with different cycling periods spanning a range from about 6 to 45 minutes. Two of the units undershoot the negative dead-band limit, which may be due to a poor choice of location for the thermometers. On the right side, the collage presents the air and water temperatures in each unit. For unit #3, the air-water thermal contact is weak, as can be seen by the water temperature continuing to rise throughout the experiment. Unit #5 has such close thermal contact between the water and air that the air cannot reach the lower temperature limit without the water being cooled more than half a degree Celsius, which takes tens of minutes.

In the first experiment, the air temperature setpoint for all model houses was 25 °C and 200 W of heat was supplied to the water heater in each one. The temperature deadband for the thermostat function was set to 3 °C. Early data for the temperature response of the systems without adjustment is shown in Fig. 11. There is a wide range of cycle durations for these settings, from about 6 to 40 minutes, and there is no evidence of phase locking.

Model houses #3 and #5 both contain 30 gallon water heaters rather than the 20 gallon tanks in the other units. In spite of their nearly identical construction, #3 had much shorter cycle durations than #5, and it also had weaker thermal contact between the air and water. The latter is seen in both the difference in temperature between the air and the water in the hydronic loop and in the fact that the water temperature had still not reached an equilibrium range in Fig. 12. Disturbing the heat exchanger to clear any air bubbles and repositioning the air thermometer succeeded in improving the thermal contact in unit #3. For unit #5, the cycle duration could be adjusted to be shorter by reducing the water flow rate through the heat exchanger, thereby attenuating the water-air coupling, as shown in Figure 13.

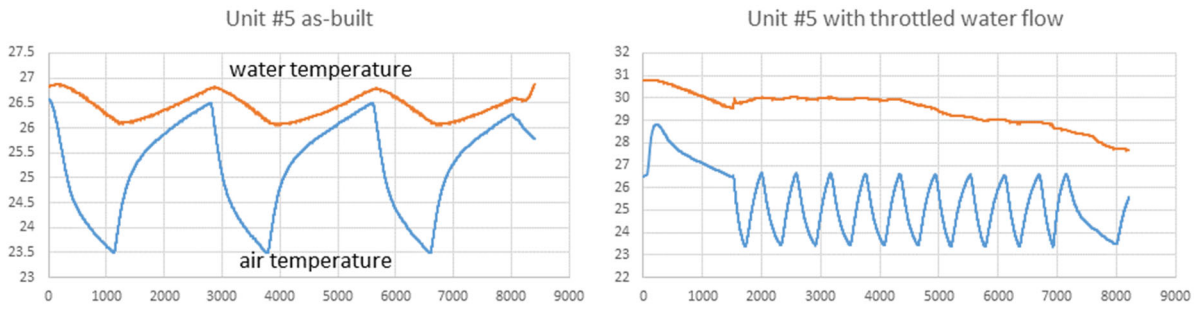

Fig. 13: Reducing the coupling between the heat capacity (water) and air results in the time-averaged water temperature running hotter. The mean heating rate remains the same, and so must the duty cycle of the AC by conservation of energy into the model house. However, the cooling power of the AC can dominate the air heating rate until larger air-water temperature differences at attained. This reduces the time for a complete thermal cycle, for a fixed heating rate. The period of cycling is about 5 times shorter for the plot on the right with throttled water flow.

UM has compiled cycle duration data from the PSI field testbed into histograms for different outside temperature ranges [Fig. 14(a)]. The cycle durations are collected from 47 real houses over several summer months, with extreme conditions for power consumption or cycle duration removed. For the experimental testbed, it is not possible to control the lab temperature, and the lab temperature does not generally approach the temperatures in the PSI data set (for houses in Austin, TX). However, the experimental conditions of temperature setpoint, temperature deadband, and internal heat load can be controlled to provide different steady-state cycle durations. These data can be compiled from different experiments to show typical cycle durations for the model homes in Fig. 14(b). The experimental testbed units have durations that span the useful range for testing controls, with most units near the mean value seen in the PSI data.

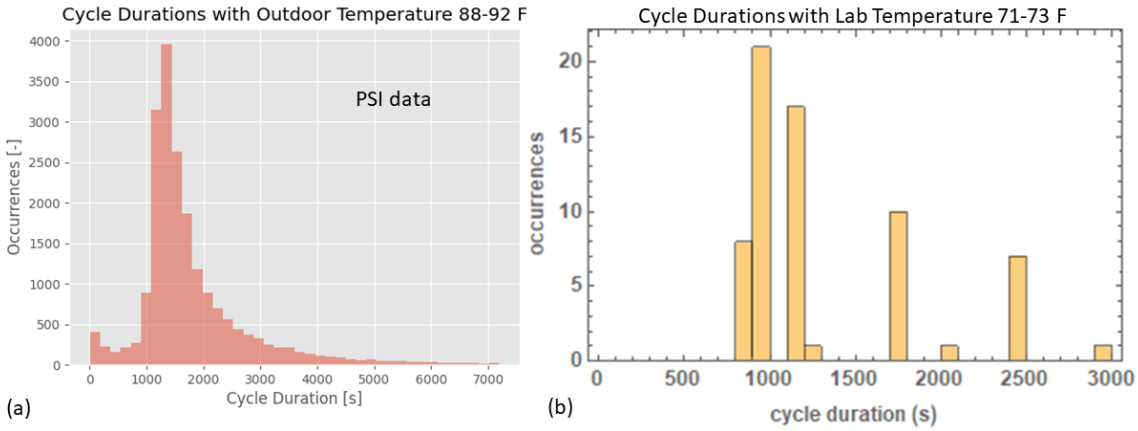

Figure 14: TCL cycle durations for (a) a compilation of PSI field test data, and (b) a sum over experiments on the experimental testbed. One unit in the experimental testbed was cycling with period ~500 s, with this peak biasing the overall distribution to lower durations. However, the cycle duration for that unit was increased by widening its temperature deadband so that it would not dominate the distribution.

The eGauge updates all power readings approximately once per second. Real power for all the TCLs in the first experiment is shown in Fig. 15. Comparison of on/off times for each AC results in duty cycles of about 30-35 % for all units except #3, for which the duty cycle is about 20%. The duty cycle for unit #3 is lower than the for the others because it is not removing all of the time-averaged heat going into the model house, due to the poor air-water thermal contact described in the Fig. 11 caption above. Had we waited long enough for the water temperature to saturate, the duty cycle would approach that of the other units in the steady state.

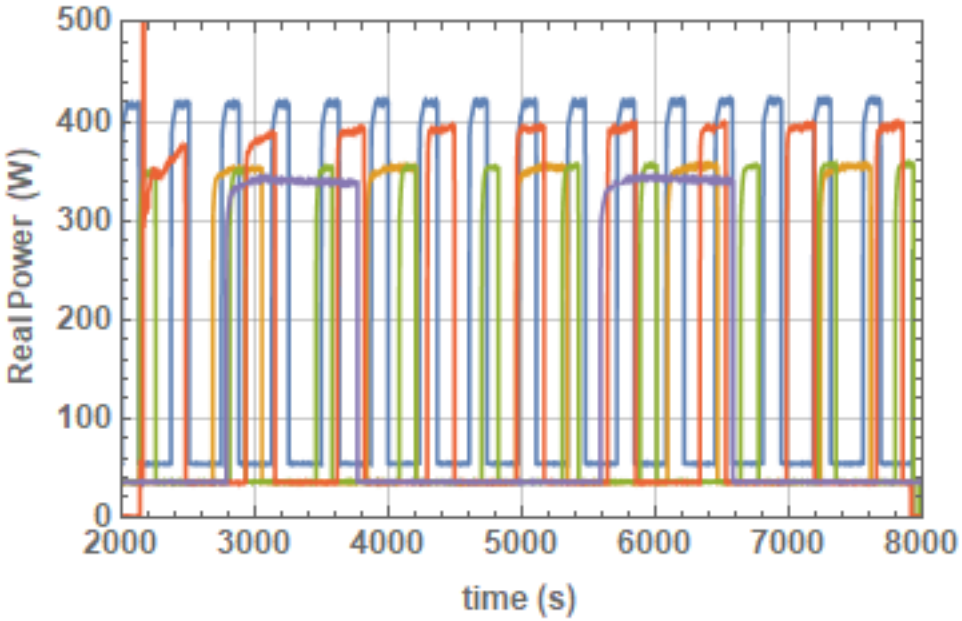

Fig. 15: Real power consumption by the 5 ACs. The highest power usage, 420 W, is for unit #1, which is a different model AC than the others. The fan for unit #1 also uses more power than the other ACs, 56 W vs. 38 W. Units #2, 3, and 5 have peak powers of about 350 W.

The base load of the fan and water pump generates about 120 W of heat into each model house. This load is not controlled and is needed to operate the hydronic loop. The baseline power consumption of the air conditioners, 38 W to run the fan on "High", does not contribute much to the heat load inside the box because only the fan's shaft enters the cold space: The motor itself is on the "hot" side, and heat generated there is rejected directly to the laboratory. In this early experiment, the additional 200 W into the water heater, controlled by the computer, led to a ~35 % duty cycle when the systems equilibrated. In that case, we can expect to vary the duty cycle of the single-speed air conditioner by adjusting the time-averaged heat input of the water heater from

120 W / 320 W * 35% = 13%,

for Qwater=0, to 70% for Qwater=520 W. This range of duty cycles covers the bulk of the distribution seen in the field test data from PSI shown in Fig. 16. In the experimental testbed, a similar *distribution* of duty cycles can be arranged within a single experiment by setting a range of different heating rates across the ensemble of model houses. To achieve higher duty cycles in any model house, we would have to increase the maximum heating rate allowed, as described in the next section.

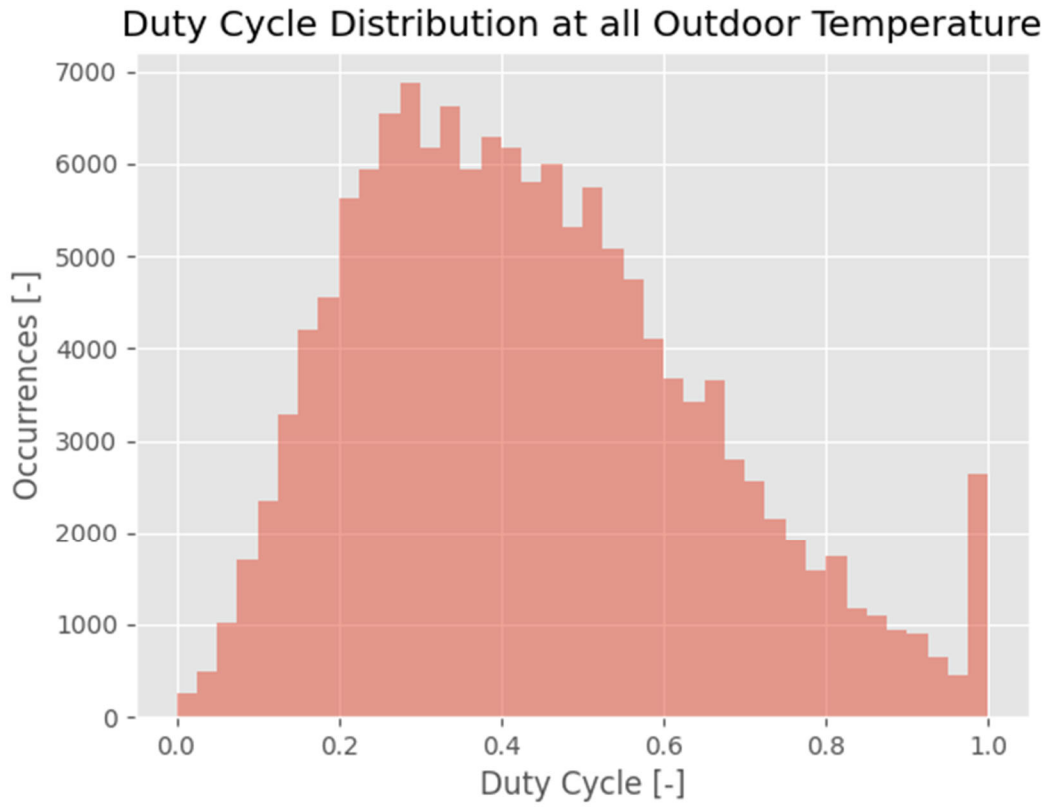

Figure 16: Distribution of duty cycle for TCLs in the PSI field test data. This data is compiled over all outdoor temperature conditions.

Typical real and apparent power (P and S) for one of the TCLs is depicted in Fig. 17. The real power reaches a plateau within about 10 seconds, and the apparent power has a transient just when the AC is turned on due to the inrush current. The eGauge records this spike in apparent power every time the AC compressor is switched on. Closer examination of the real power when switched on shows a smaller spike there as well, but often the real power spike is reduced to a point of inflection of the curve. In general we take the first point of the ON-state to be the real part of the inrush power transient and compare this to the steady-state power with the AC on, forming an inrush power to steady-state power ratio. The experimental testbed values for this ratio for one experiment are compared to the distribution derived from PSI's field test data in Figure 18. The agreement is not perfect, but most of the occurrences for the experimental testbed occur at a ratio of nearly 60%, which is within the PSI data's distribution. The air conditioner units in each study are quite different, with the PSI units using about eight times the power of the small window air conditioners of the experimental testbed. Also, the power factor (P/S) of the PSI units is about 0.875 compared to 0.988 for the ACs in the experimental testbed.

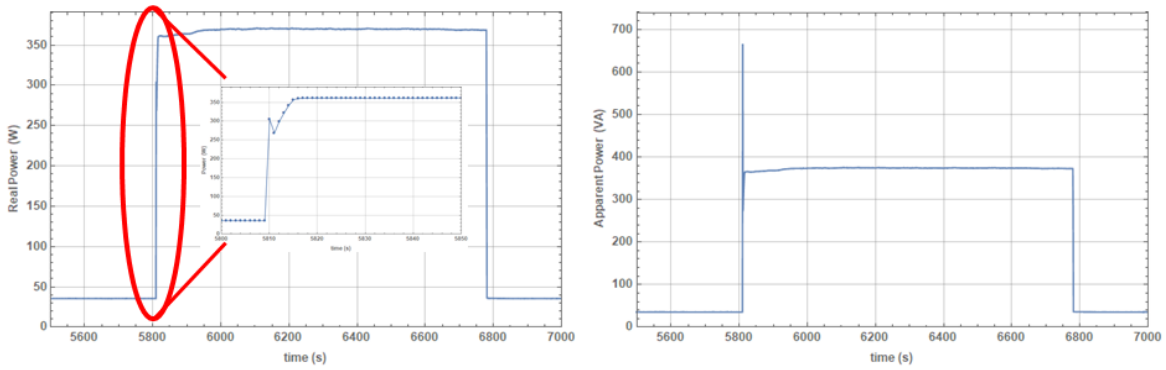

Figure 17: Real (left) and apparent (right) power during a single on/off cycle of the AC compressor. The inset shows the leading edge of the ON-transient, which has a discontinuity due to the inrush current at startup of the compressor.

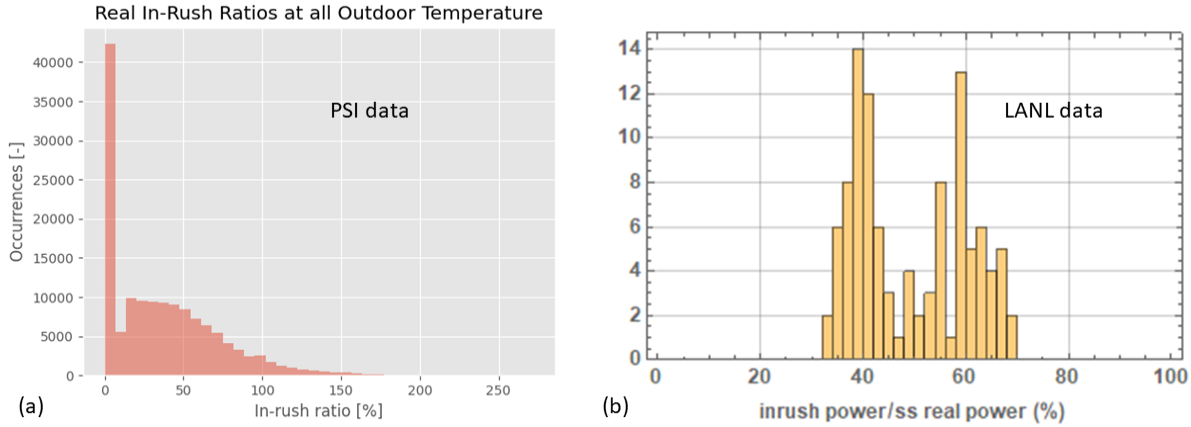

Figure 18: (a) The ratio of real power during the inrush transient to the steady-state real power is computed for all ON-transitions in the PSI data to populate a histogram. (b) The same ratio of powers was computed for a single experiment in the experimental testbed. Most of the occurrences at 40% are from Unit #1, which had a much shorter cycle duration than the other TCLs.

Although the lab temperature cannot easily be controlled, it may drift 2-4 °C over the course of an experiment. The drift depends on outside temperature and solar irradiance of the building, on co-located activities that may open the roll-up doors to the outside, and on the local heating by the experiment itself. The real, steady-state power consumption of the TCLs is seen to depend on the lab temperature, presumably because of the temperature changes imposed on the condenser-side heat exchangers of the ACs by lab temperature. In other words, the real power and lab temperature are correlated, as shown in Fig. 19. For the LANL testbed, each "occurrence" is the calculated correlation for one of the five initial model houses in a particular experiment. The correlations are in reasonable agreement with the correlations seen in the PSI data. The power dependence with temperature is so clear on a linear plot that one can extract a local derivative dependence of

$(dP/P)/dT = 6/380 = 1.6\%/°C$,

compared with a trend of approximately 1.8%/°C for the PSI data and approximately 2%/°C for the residential AC unit in Fig. 6 of the WECC load modeling research project of Ref. [6].

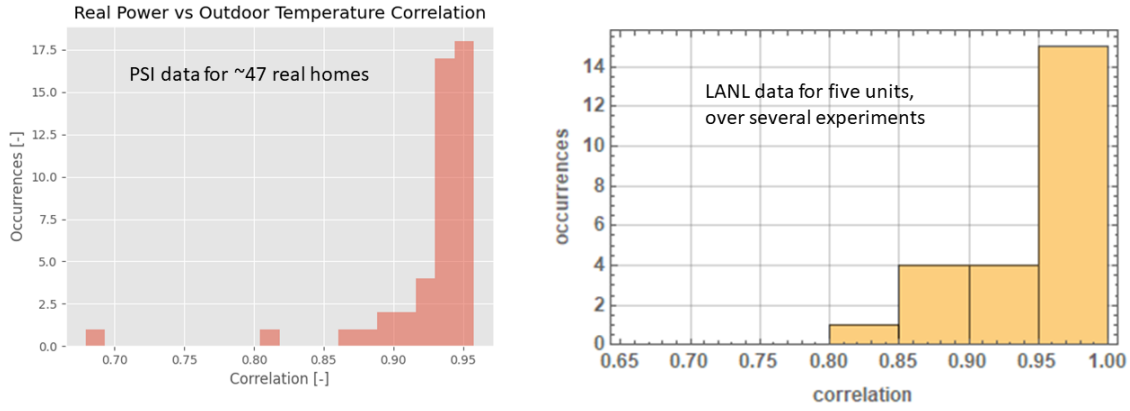

Figure 19: Correlation between outside temperature and peak power. (a) PSI field test data. (b) LANL testbed experiments on several days.

The relative phasing of ON/OFF transitions between pairs of TCLs was investigated to look for evidence of phase-locking among units. That is, the timing of ON-transitions for the longer-cycle-duration TCL was used as the zero of time, and the time to the next ON-transition of the other TCL was divided by the longer cycle duration:

$\Delta\theta = (t_2 - t_1) / (\text{cycle duration 1}) * 360°$

These phases were plotted against successive transitions of the slower-cycling TCL as shown in Fig. 20. Although some TCLs had periods that were close to rational fractions of one another, there were no clear cases of coupled oscillation.

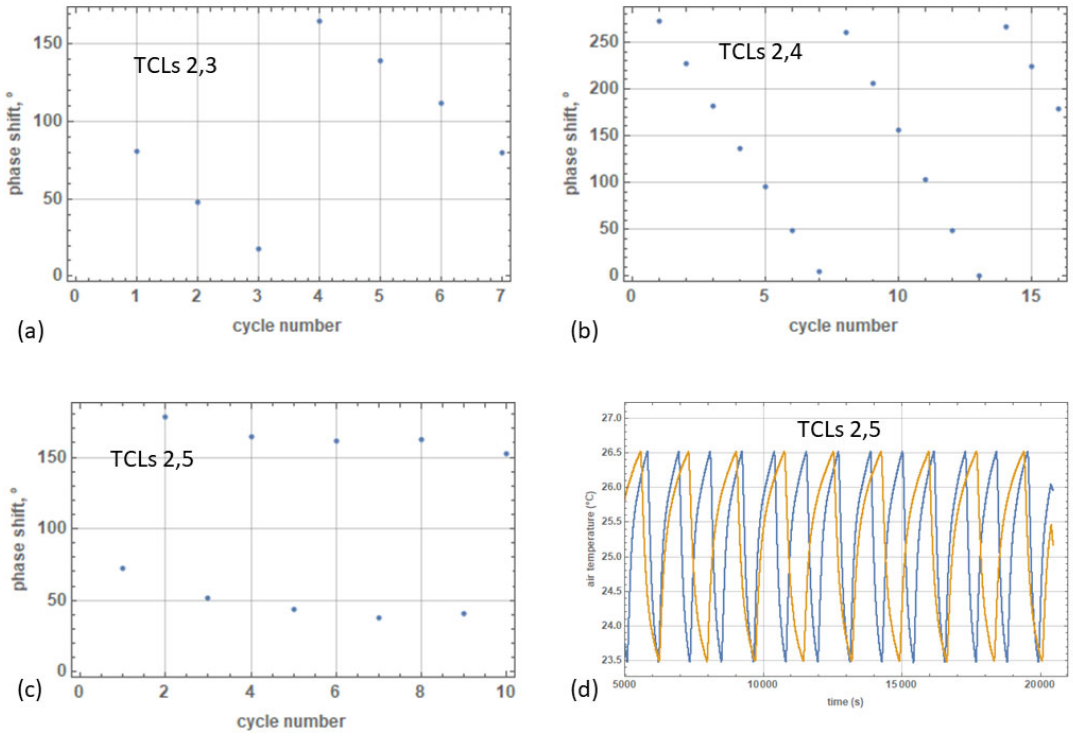

Figure 20: In (a) and (b), there is roughly constant phase drift between the ON-transitions of the pairs of TCLs. In (c), the relative phasing seems to alternate, but this only indicates that their periods are close to a ratio of small integers and not necessarily phase-locked.

Finally, experiments were carried out that tested the forced ON/OFF switching of the five initial TCLs together. Alternating ON and OFF requests were sent to the "software follower" subroutine every 180 s, and the follower would determine whether to set the newly requested state and would send a response code describing the resulting action for each request. The air conditioner specifications require that the compressor not be turned back ON within 180 s of having been turned OFF, in order that the compressor not stall upon restart. Therefore, it would not be meaningful to make alternating state requests faster than once every 180 s. There is no minimum time for the compressor to be ON, so the switching time constraint is asymmetrical:

Lock-out time until next ON-state = 180 s, and

Lock-out time until next OFF-state = 0 s.

The temperature trajectories for two experiments at different internal heating rates are shown in Fig. 21. When the TCL follows all the ON/OFF requests, it has a 50% duty cycle. The heating rate chosen in (a) of the figure is lower than half the AC's cooling rate, so the model house cools

down to the lower temperature limit. For (b), the heating rate is much higher, so the model house warms up to the upper temperature limit.

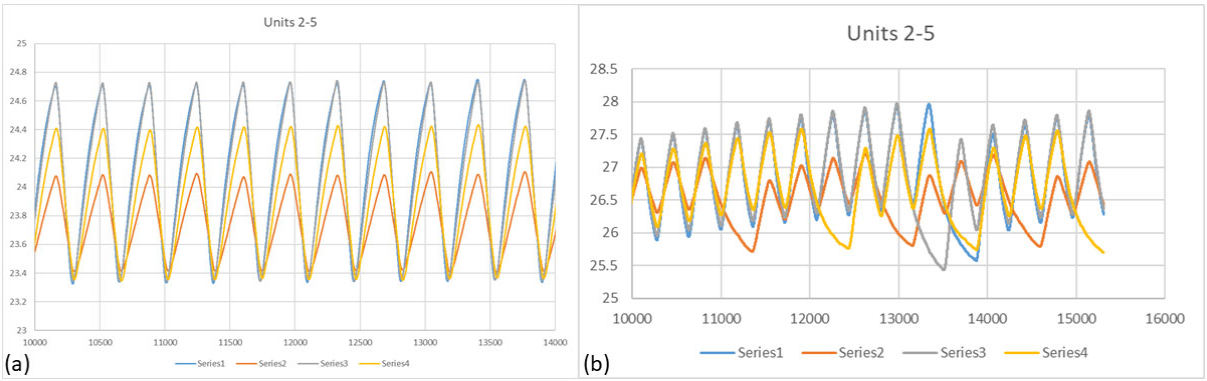

Figure 21: Experiments with forced switching every 180 s, for two different applied heat loads in the model houses. The model house temperatures have settled to the lower or upper temperature deadband limits [23.5 °C or 26.5 °C for (a) and (b), respectively] depending on whether the AC's cooling rate at 50% duty cycles is greater or less than the applied heating rate. (a) The TCLs all turn off when the temperatures reach the lower edge of the deadband around 25 C. Because there is no minimum time for the compressor to be ON, the OFF request is always accepted. (b) At the upper deadband limit, the compressor may need to remain OFF even though the temperature is above the limit, due to having switched OFF too recently. In that case, the TCL must wait until the 180 s minimum OFF-time (the short-cycling limit) is complete before the thermostat can switch it ON again, and the temperature inside the model home keeps increasing until that occurs. When the compressor is finally allowed to turn ON, it may not be able to cool all the way back into the deadband before the next OFF request arrives. Then the follower rejects the OFF request, and the TCL will cool into the deadband for another 360 s or until the lower temperature limit is reached. Thus, some requests are rejected in (b) but the ACs always switch together in (a).

The model houses in the experimental testbed generally execute the same cycles over and over during open-loop operation, due to the constant internal conditions and nearly constant conditions in the lab space. In contrast, the individual houses in the PSI field testbed have time-varying internal heat loads and widely varying external conditions of temperature, humidity, and solar irradiance. While it is not possible to use a single experiment with the five trial TCLs to fully populate the various distributions from the PSI data set, the tests above show that the trial TCLs mostly fall within those distributions and can be tuned to different positions within them. Some characteristics of the PSI loads cannot be exactly matched due to scale differences, but for the most significant (e.g., cycle duration) the heterogeneity can be more closely copied.

2.5 Design Reviewed and Revised Based on Validation Results

The SOPO describes Task 2.5 (design reviewed and revised based on validation results) as follows:

> *Deviations of the trial TCLs from the PSI data investigated. Insulation, positioning, heat load range, and/or ventilation changed as necessary to minimize deviations from the PSI data.*

This task is complete. As described in the previous section, the cycle durations depend strongly on the air-water thermal coupling, and identically constructed model houses may therefore have large differences in these durations. Fig. 12 above shows tuning of the cycle duration to lower values by throttling the water flow in the hydronic loop with a valve; and short durations can sometimes be lengthened by taking measures to eliminate vapor bubbles or otherwise improve thermal contact to the heated water. If the latter measures are unsuccessful, increasing the width of the temperature deadband generally lengthens the cycle duration, because the increased air-water temperature gradient drives greater round-trip heat flow per cycle.

Cycle duration is affected by changing the heating rate to the water tank, but the main effect is on the duty cycle. Duty cycle varies approximately linearly with applied heat because the single-speed air conditioner only transfers heat at a constant rate. The initial five heater control boxes incorporated a 60 V, 625 VA power transformer to limit the power into the 2000 W-rated water heaters to 500 W and provide finer control over the heating rate. The resulting duty cycle was seen to be between about 13% and 70%, but it may be useful in experiments to reach as much as 90% to better match the distribution of field test data in Fig. 15. To increase the maximum rate to 889 W, the new heater control boxes will use an 80 V, 1500 VA toroid. This also required an increase in wire thickness in the design to carry the higher current.

Other comparisons with the PSI data are more difficult to control, namely the inrush power to steady-state power ratio and the power factor of the unit TCLs. The only possible adjustment would be in choosing a different model air conditioner, and the selection of these is limited. The larger-capacity models require more power than the lab can supply for an ensemble of twenty units. However, these properties may not need to be in exact agreement for the purpose of testing controls. In summary, we expect that there is sufficient tuning flexibility to minimize deviations with PSI data.

2.6 Remaining 15 TCLs Fabricated

The SOPO describes Task 2.6 (remaining 15 TCLs fabricated) as follows:

Remaining 15 TCLs fabricated.

Specifically, final sketches from trial TCLs reviewed/revised and supplies for the remaining TCLs purchased. TCLs fabricated, transported, and installed in the lab space. Data acquisition system expanded to capture the extra channels. Instruments and sensors calibrated as necessary. Hooks for the simulation testbed to interact with experimental testbed and data acquisition system designed and implemented.

This task is complete. All voltage-to-current converter circuits for controlling the water heater in each model house have been fabricated and installed. The boards were calibrated with shunt resistors to provide a span of 4-20 mA for an input of 1-5 V from the control computer.

The Watlow SCR does not present the control current circuit with a purely resistive impedance, so one of the heater control circuits was tested with a simple incandescent bulb and with the eGauge power meter. This test verified the functioning of the SCR with the current control interface box and the linearity of the power control as measured by the meter.

The interfacing with the enhanced simulation testbed and control algorithms, running as a separate process on the same computer, can be implemented by message passing through TCP ports. LabView has built-in functions for TCP, allowing simple interprocess communications over the loopback network with the details of message packetization handled automatically. We may choose a pair of ports for this bidirectional communication. Examples of bidirectional communications between LabView and MATLAB were located, and a LabView-to-LabView example was coded and tested for proof-of-principle and for testing the speed of passing data arrays as message strings in JavaScript Object Notation (JSON) format. The controller for "forced-switching" experiments, where the controller requests ON or OFF transitions at fixed intervals, was implemented in this way as a separate program running alongside the data acquisition program.

Experience from building the initial five units helped identify improvements that reduced construction time in the remaining set. For example, the ¾" NPT fittings throughout had proved difficult to seal reliably and occasionally had hairline cracks that leaked. In the remaining 15 houses, NIBCO press fittings were used in some locations, reducing the time and labor in assembly and leak checking in exchange for a small increase in cost per fitting. Similarly, Sharkbite fittings were used for the PEX tubing junctions, reducing assembly time.

The 15 final TCLs were completed and installed. Design modifications to the hydronic loop saved much assembly time. An additional technologist was brought in to assemble the 15 power control chassis. The control lines were connected at each unit and routed back to the data acquisition system. Thermistors were installed, and the wiring for voltage readout completed. Figure 22 shows the array of TCLs in place just before wiring.

Figure 22: This shows the testbed just before final connections. All 20 units are functioning.

2.7 Comprehensive Testing and Validation Conducted

The SOPO describes Task 2.7 (comprehensive testing and validation conducted) as follows:

> *The full experimental testbed comprehensively tested and validated against the PSI data.*
>
> *Specifically, data from experiments on the complete testbed compared with the data from the trial TCLs. Data examined for phase locking or coupling. Performance assessment completed, specifically, whether (scaled) TCL real and reactive power consumption during normal operation in the experimental testbed is within 5% RMSE error with respect to the PSI data*.

This task is complete. The benchmarking experiments with 20 TCLs are finished. Experiments were performed at several temperature set points, with different temperature deadbands, and with different internal heating rates. The data from the complete testbed is qualitatively the same as that from the initial five TCLs with similar distributions of cycle durations and duty cycles. No coupling or phase-locking between units was observed in steady-state operation.

Validation consisted of the following steps.

1. Used the PSI data to plot histograms of real TCLs:
    a. Duty cycles
    b. Periods of cycles
    c. In-rush power to steady-state power ratios
    d. Variation in power consumption when TCLs are ON

and compared the histograms to the values obtained from the complete set of TCLs.

2. Checked if the variation in power consumption when the complete set of TCLs are ON is correlated with
    a. Voltage
    b. Ambient temperature

    and compared the correlations to those obtained using the PSI data (Milestone 2).
3. Checked for expected behavior when trial TCLs are forcibly turned ON/OFF by a controller.

The results of these tests on the five initial model houses were described in Sec. 2.4 above. Adjustments to TCL characteristics may be made by methods such as

a. Changing temperature setpoints (straightforward)
b. Changing heat supply to water tank (straightforward)
c. Changing TCL dead-bands (straightforward)
d. Changing ambient temperatures (harder, requires construction of a foam shell)
e. Throttling water flow and diverting airflow (harder)

Several of these options were tested as described in Sections 2.4 and 2.5 above.

Step 1a.

For validation step 1a above, a graph of duty cycle from the PSI data was given in Fig. 15, back in Sec. 2.4. Here we show duty cycle data from the full testbed at typical conditions, in Fig. 23. The PSI data, corresponding to all outdoor temperatures, shows a broad peak with its maximum around 30% duty cycle. In contrast, Fig. 23 shows distributions corresponding to a single heating rate for all TCLs resulting in narrower distributions with different peaks. A broad distribution of duty cycle similar to PSI may be obtained by choosing a distribution of heat inputs to the various TCLs instead of using a single heating rate for all. The modified power modules (described in Sec. 2.5) source enough heat in the model houses to exceed the cooling capacity of the air conditioners at a duty cycle of 100%. Therefore, a reduced-sample approximation of the duty-cycle plot of Fig. 15 can be constructed by selection of heating rate for each model house in the ensemble. The main limitation is there being a minimum duty cycle of 15% to remove the fixed internal heating of the fan and water pump. Heat leaks to the lab's ambient temperature will also contribute to duty cycle for each model house.

In summary, we can achieve nearly the full range of duty cycles seen in the PSI data, and we could tune the heating rates to achieve a distribution similar to the PSI distribution. Therefore, we have validated the duty cycles.

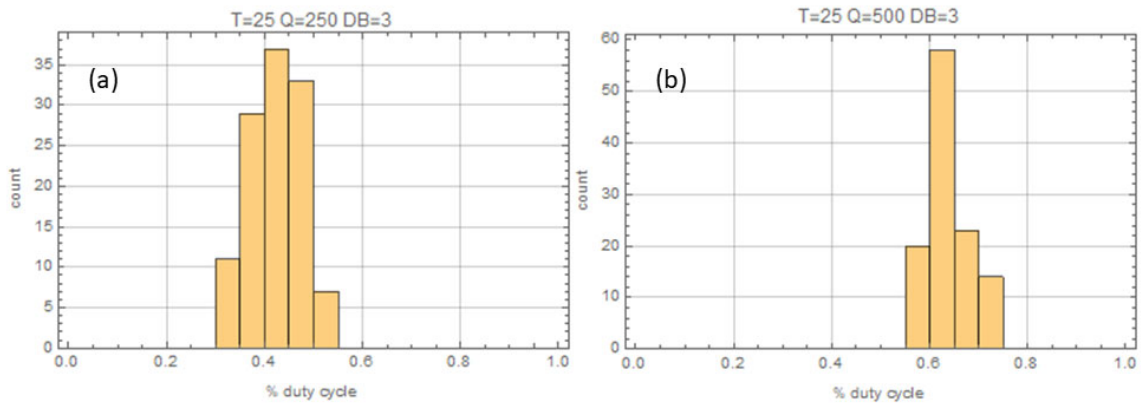

Figure 23: The experiments in (a) and (b) have the same setpoint of T=25 °C, but Q=500 W of heat is injected into the model houses for (b) vs. Q=250 W for (a).

Step 1b.

For validation step 1b, cycle durations for the TCLs are shown for four experiments in Figure 24, to be compared with the distribution from the PSI data shown in Fig. 14(a). The clearest feature from these experiments is that decreasing the temperature deadband width in 24(d) shortens the cycle durations. For the other experiments, the distribution of cycle duration is peaked around 1000 to 1200 s which is comparable to the cycle durations seen by PSI. It may be useful to reduce some of the longer cycle times for testbed experiments, because the TCLs will otherwise execute only 3-4 cycles over a 2-3 hour experiment. The faster-cycling TCLs may be adjusted to longer cycle durations by increasing the width of temperature deadbands for their thermostats. Therefore, there is sufficient flexibility to adjust these TCLs to closely match the cycle-duration distribution from PSI--without altering duty cycle--if desired.

Summarizing this step, we can tune the distribution of cycle durations in the ensemble to agree with the PSI data using the deadband width for each TCL's thermostat. This validates the cycle duration performance of the experimental testbed.

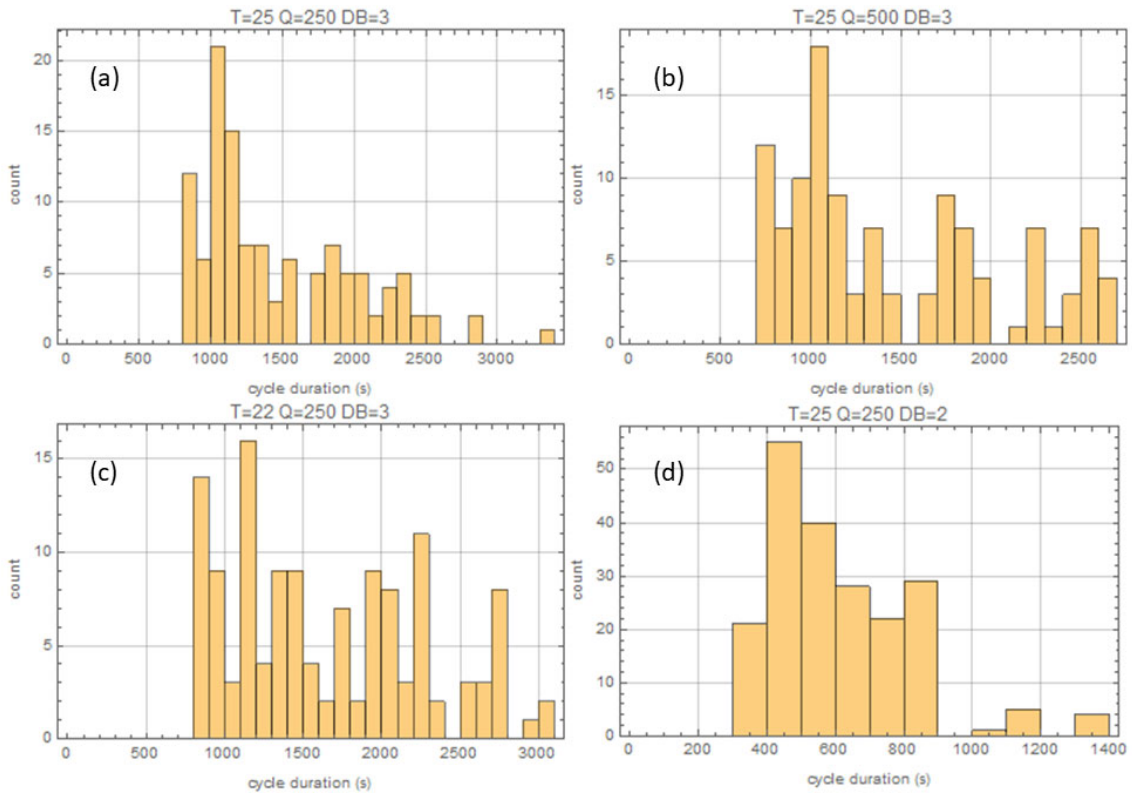

Figure 24: Cycle durations for the testbed with no physical adjustments, for four different experimental conditions. (T=setpoint in °C, Q=heat injection in W, DB=dead-band width in °C)

Step 1c.

Validation step 1c concerns the in-rush power draw of the air conditioners. Real and reactive in-rush power are compared against the steady-state values of each phase, respectively, in Fig. 25. In the histograms, the in-rush is considered as a fractional deviation above the steady-state value:

% deviation = ([real or reactive in-rush power]/[real or reactive steady-state power] – 1) *100 %

In the LANL experiment, the first power measurement after turning ON is considered the in-rush, because the compressor is motionless for the beginning of the measurement window. Faster measurements (at 1 kHz) of the current show that the in-rush always spans about 10 cycles at 30 A before settling down to the normal 3-4 A currents at steady state. Negative deviations in the real power were measured (i.e., the real inrush spike was lower than the steady-state real power), but these were taken as zero deviation for the purpose of comparing with the histogram of real PSI data. Since most of the LANL real in-rush power measurements were less than the steady-state power, the zero-power bin dominates the distribution.

The in-rush deviation distributions are not similar between PSI and LANL data. This is not completely surprising, because the air conditioners types are different. The PSI air conditioners are higher power and use reciprocating piston compressors. The LANL air conditioners, though, are small window-mount units that operate from low-inertia, rotary compressors. In the Milestone 2 report, Fig. 6, the in-rush power spikes of the PSI air conditioners are essentially random and likely depend on the phase at which the compressor stopped. The LANL air conditioners, in contrast, have similar spikes throughout a given experiment. There may be some variation or drift in the value of the spikes recorded by the eGauge during a given experiment and larger differences from run to run, but these values to not appear random. Rather, the measurements may drift slowly due to differences in the time base of the eGauge with respect to the LabView control computer.

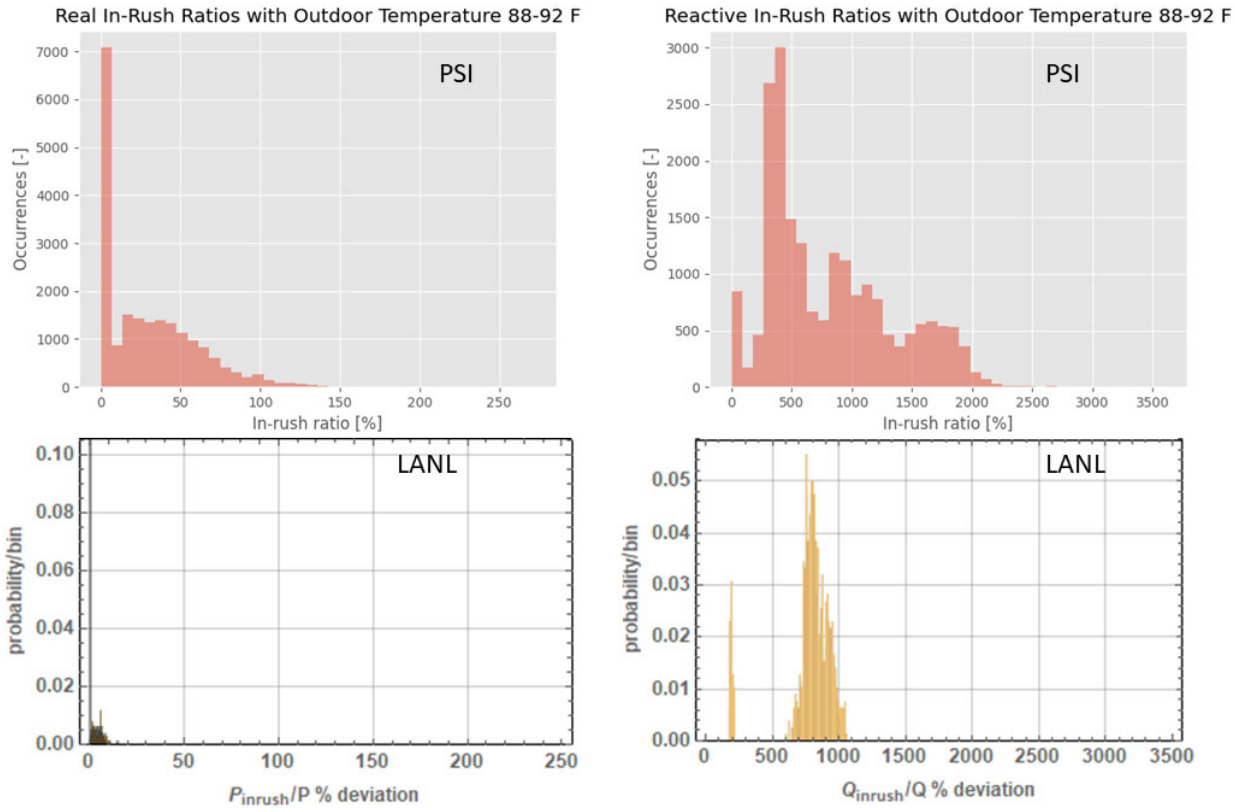

Figure 25: Histograms of real and reactive in-rush power vs. steady-state power, for both PSI field measurements and the LANL experimental testbed. Zero % represents an in-rush spike that was not detected (PSI) or did not exceed the steady-state power in that phase (LANL). LANL data is represented as probability density in each bin; and for the real in-rush, about 90% of the measurements were below the steady-state real power value. The peak in the reactive in-rush for PSI data occurs at about half the value of the peak in the LANL data. However, the TCLS of the PSI data generally have lower power factors than the LANL data. For example, Fig. 6 in the Milestone 2 report shows an AC with power factor ~0.9 vs. 0.99 for the LANL testbed ACs. That is, the reactive in-rush power deviations seem much higher in the LANL data because the steady-

state reactive powers in the denominator are proportionally lower. The TCLs in the PSI data set have mean power factors ranging from 0.87 to 0.99.

Summarizing task 1c, the in-rush characteristics are not identical for the PSI and LANL data. There is less heterogeneity within the LANL data, and the LANL distributions are bounded by the PSI distributions. In-rush is not an adjustable characteristic, depending on the start-up of the AC compressor motors. Therefore, it will be more useful to consider the in-rush behavior of these ACs in the simulation testbed modeling, as the outcome of this validation step. The impact of different in-rush characteristics can be further explored in Milestone 6, where the simulation testbed can "add" in-rush current to the simulated distribution network buses where the experimental TCLs are "located" to investigate the impact of in-rush on power flows, voltages, and so on.

Step 1d.

For validation step 1d, the variation in the power consumption of each air conditioner was evaluated in the ON state. For each ON state, the minimum and maximum power were found after discarding the first and last 5 seconds of data to ensure removal of transients. Then the fractional variation is calculated from

Fractional change = (P_max – P_min)/P_max

The fractional power variations were combined for all air conditioners in each testbed, and results are plotted in Fig. 26. For the LANL testbed, the mean power variation is 10 +/- 3%, whereas for the PSI testbed it is 8.7 +/- 9.3%. The large variation in the PSI data is contributed primarily by two air conditioners with ~35% deviations on some cycles.

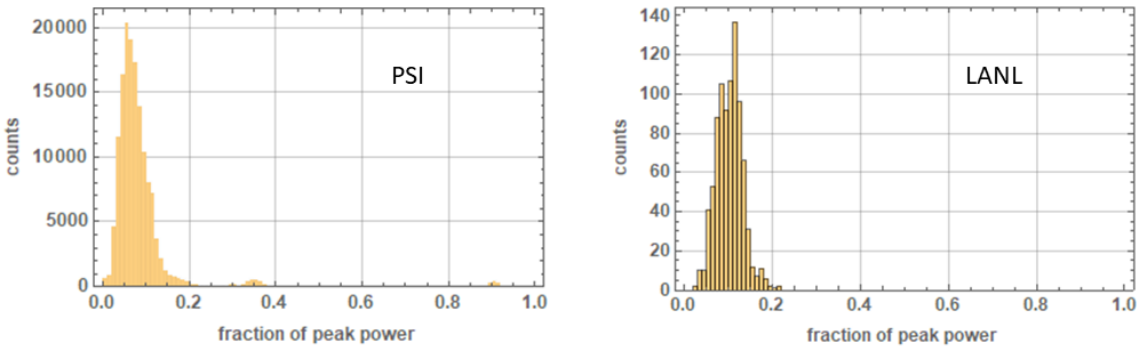

Figure 26: Power variations scaled to peak power for PSI and LANL testbed populations.

Summarizing step 1d, the power variation distributions of the two testbeds agree within their uncertainties. There are no adjustable parameters to improve the agreement. Therefore, this validation step is complete.

Step 2.

To approach validation steps 2a and 2b, data sets from multiple experiments on the full testbed were joined together to calculate the correlations of real power P and reactive power Q vs. line voltages and lab temperature. Before showing histograms of the calculated correlations for each AC, we show typical scatter plots upon which these correlations are based, first for the PSI data in Fig. 27 and then for the LANL data in Fig. 28. These correlations are for a single TCL in either figure. There are 47 TCLs monitored over an entire summer in the PSI data set, and there are 20 TCLs measured over 8 experiments in the LANL experimental testbed.

From the scatter plots, one sees that the correlations of P and Q vs. voltage are quite weak in both testbeds. The real power appears to be negatively correlated with voltage in the PSI data, but it was recognized in the M2 report that this apparent correlation was actually a result of variations in the ambient temperature. When the outside temperature was high, more air conditioners were turned on in the field testbed, and this increased loading of the feeder pulled down the line voltage. Voltage was negatively correlated with ambient temperature in the field testbed. In the experimental testbed, though, all ACs are working in all experiments irrespective of the laboratory temperature, so voltage and ambient temperature are mutually independent variables and therefore uncorrelated.

Turning to the scatter plots of P and Q versus ambient temperature for validation step 2b, P and T are strongly correlated in both the LANL and PSI data. Linear fits are made to the data from ACs in each testbed to model the fractional change in real power consumption with outdoor or lab temperature. The mean fractional change in power is about 0.80 %/°F for the TCLs in the LANL testbed and 1.18 %/°F for the PSI testbed. While these mean values differ by more than 5%, the scatter among the LANL units is 6.3% relative and among the PSI units it is 13% relative. The LANL TCLs' power variations fall at the lower end of the PSI TCLs's distribution. The temperature dependence of power consumption is thus in fair but not perfect agreement between the testbeds. The air conditioners in the testbeds are of different models and capacities, and there are no adjustable parameters in the comparison, so perfect agreement could not be expected.

For the reactive power Q, there is at most a weak correlation with temperature in both data sets.

Thus, we find that the correlations of real and reactive power with voltage and temperature are in reasonable agreement between the LANL and PSI testbeds, completing steps 2a and 2b of the validation. The power vs. temperature correlation histograms are shown for the PSI and LANL TCL populations, respectively, in Fig. 29.

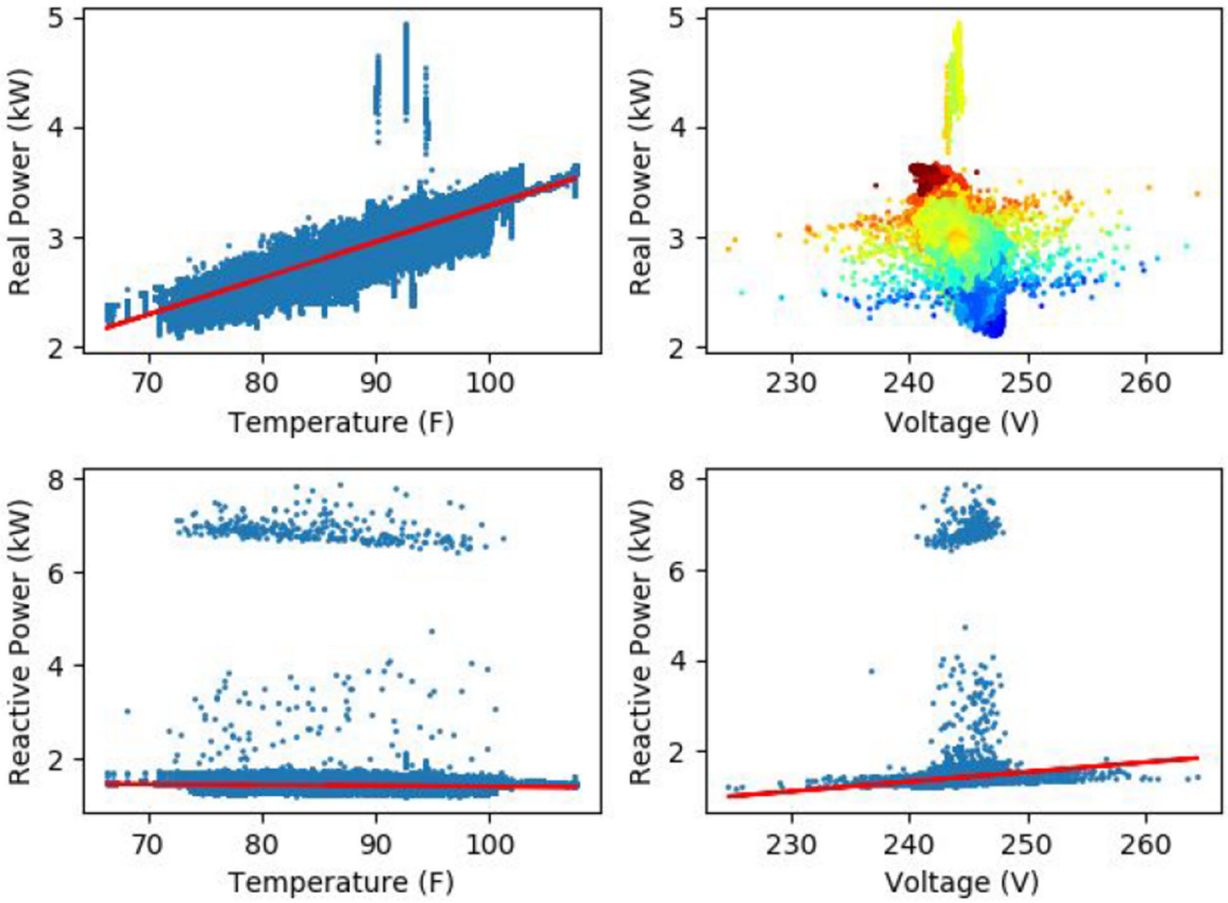

Figure 27: Single-AC scatter plots of power vs. temperature and voltage from the PSI field testbed, borrowed from Fig. 8 of the Milestone 2 report. The voltage is thought to be dependent on outside temperature in this field test data because of the increased load on the grid from air conditioning, so that true the dependence of an individual air conditioner on voltage is likely small.

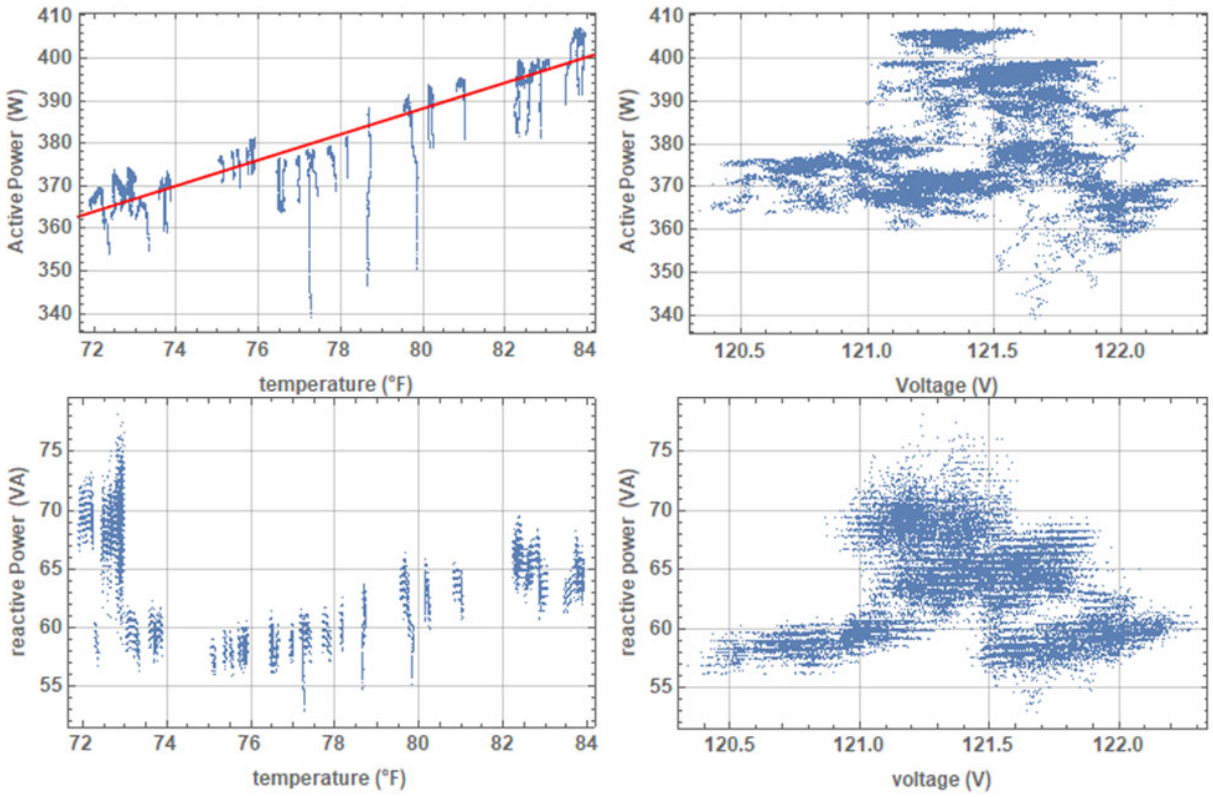

Figure 28: Single-AC scatter plots from the LANL experiment. The line voltage is not itself correlated with the lab temperature (not shown), although voltage does depend on the instantaneous power draw of the whole experiment. Only the power vs. lab temperature plot has a clear trend, shown in red.

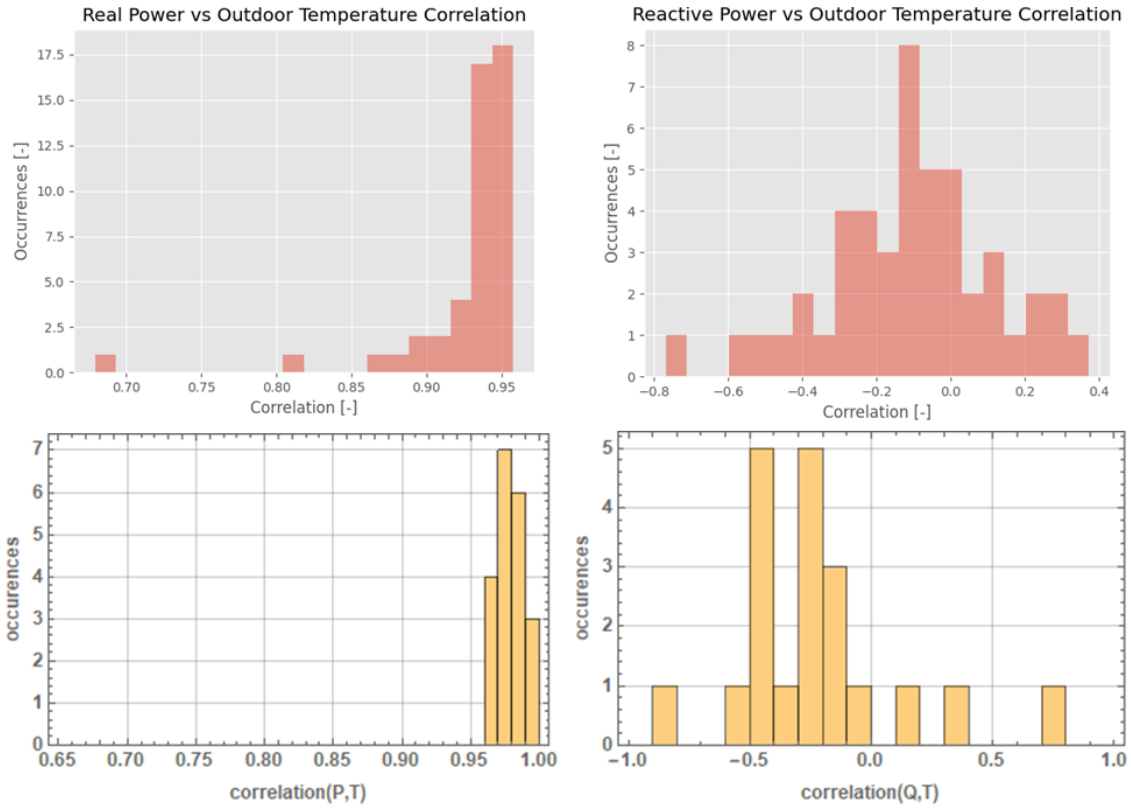

Figure 29: The experimental testbed's power data is correlated against temperature in the bottom two graphs, for real and reactive power. The top two graphs are from field test data by PSI. Note the difference in x-axis scale of the histograms for reactive power in the plots on the right side.

Step 3.

For validation step 3, we have run forced-switching experiments, which intentionally synchronize the TCLs, on the full testbed at two different periods and two different heating rates so far. These are the same as the experiments described in Sec. 2.4 carried out on the five trial TCLs. Data for the full testbed switching every 180 s is shown in Fig. 30, for comparison against Fig. 21. Again, some of the houses exceed the upper temperature limit T+ because the lockout period prevents turning ON the compressor as requested by the thermostat. The response codes resulting from controller state requests are shown in Fig. 31.

Repeating the experiments with switching intervals of 360 s, in Fig. 32, the resulting temperature behavior is about the same for low internal heating rates (250 W). For the higher heating rate (500 W), though, the model houses no longer reach temperatures above T+, because the longer ON/OFF times do not allow the houses to reach T+ before the lockout period of 180 s expires. For even longer switching intervals, the TCLs with shorter cycle durations may change state several times before the next controller request, leading to instances where the

compressor is locked out with temperature near T+. In that case, some TCLs will seem to exceed T+ randomly.

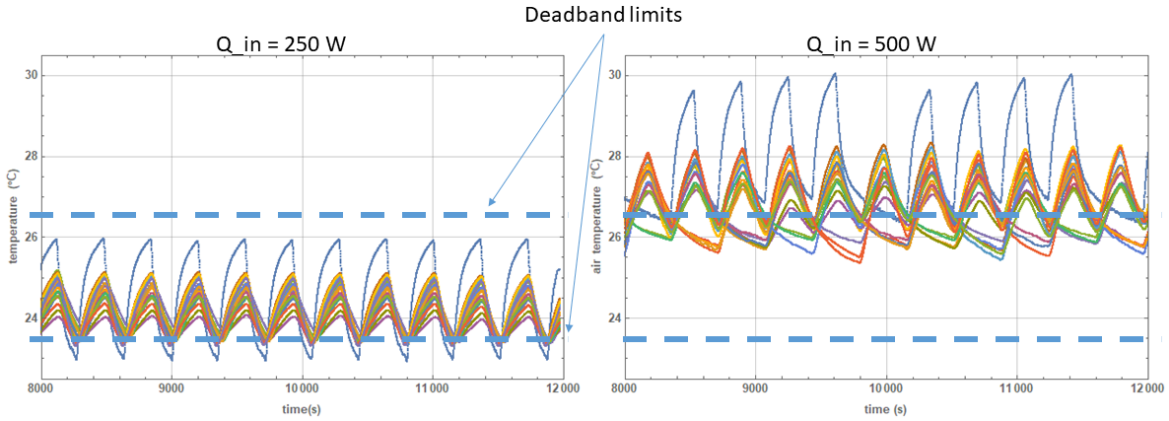

Figure 30: Experiments with compressor switching requests at regular intervals of 180 s, which intentionally synchronize the TCLs. One of the TCLs is operated with a wider deadband than all the others (4 °C vs. 3 °C). For the plot on the left, the internal heating would require less than 50% duty cycle for the AC to maintain steady temperature. The TCLs do not all reach the lower temperature limit at the same time to enter the OFF state, but they do all turn ON at the same time when requested by the controller. For the plot on the right, the internal heating rate is much higher. The model houses do not cool deeply into the deadband before requests to turn OFF the ACs arrive, and then the compressors are locked out for 180 s. As a result, the houses warm up through the upper temperature limit. The TCLs turn ON and OFF at the same time, as requested by the controller, except when ON but still above T+.

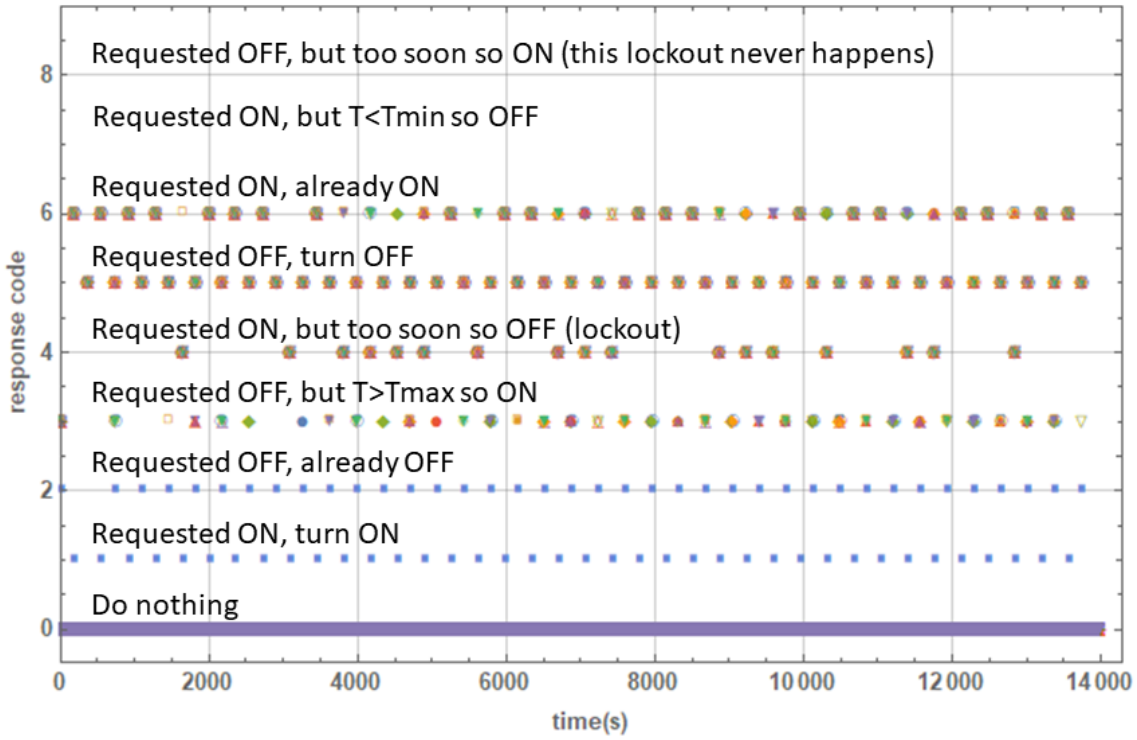

Fig. 31: Reponses from the software follower to the controller's requests, with descriptions. In this experiment the injected heat was 500 W, which leads to steady-state duty cycles of 65-75% in the model houses, given the fixed cooling rate of the AC when ON. The time between alternating ON/OFF requests was 180s. In that case, the temperatures reach the upper limit T+, but the switching interval is too short for most TCLs to reach the lower limit T- before the controller turns OFF the compressor. Therefore, the time-averaged temperatures of the TCLs are generally above the setpoint. Response code 7 (Requested ON but T<Tmin so OFF) is not used here because the temperatures never reach T-; and code 8 (Requested OFF but too soon so ON) is never used at all because there is no minimum time until one can switch OFF.

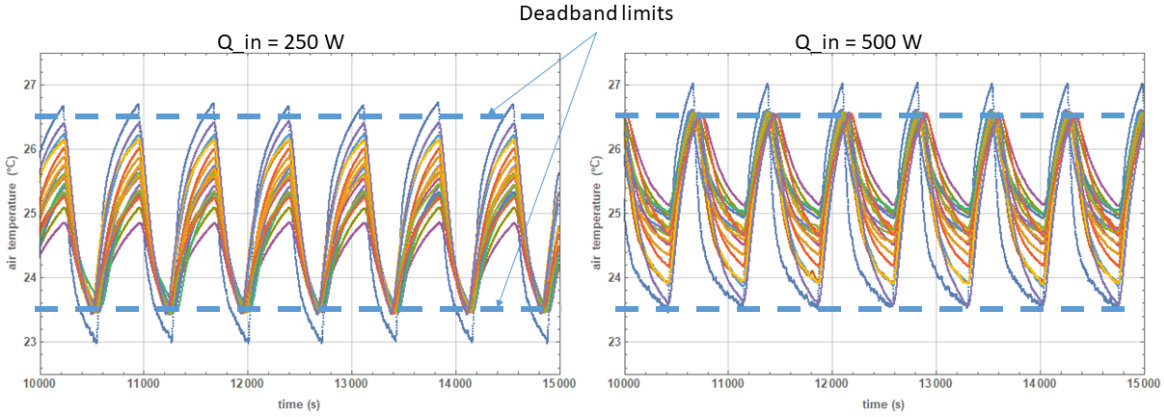

Figure 32: Forced switching experiment with 360 s intervals between commands, and with the same internal heating rates as in Fig. 29. The left hand plot is similar to that in Fig. 29. For the

higher heating rate of the right hand plot, the interval for cooling is long enough that the lockout expires before the model houses reach T+, and the temperatures stay within the deadband.

In validation step 3, we find that the behavior of the system under periodic forced switching can be understood in terms of the asymmetric lockout-time condition and the applied internal heating rate. Although the duty cycle implied by forced switching is 50%, the software thermostat and follower together keep each system within or close to its temperature deadband. The thermostat can change the ON/OFF state *before* the forced-switching request is issued, maintaining a time-averaged duty cycle equal to the steady-state duty cycle that may be greater or less than 50%. For switching times approaching the lockout time, the AC sometimes cannot be turned back ON at the T+ deadband temperature, but the follower rejects OFF-requests as needed to ensure that the model house always returns to temperatures below T+.

Summary

Validation was completed using the series of steps described in the beginning of this section. In sum, we find that the behavior of the experimental testbed agrees, or can be adjusted to agree, with that of the PSI field testbed within experimental error for most parameters. For example, duty cycles can be continuously tuned in the experiment between 15-100% for each TCL, such that one may arbitrarily match the mean and standard deviation of the PSI data. Correlations between power consumption and temperature or voltage were found to be in good agreement. Some characteristics, like real and reactive in-rush powers, do not immediately agree with the distributions from PSI probably due to fundamental differences in the functioning of these air conditioner types. In-rush characteristics in particular cannot be adjusted. However, taking into account the difference in power factors for experiment vs. field data, even in-rush power appears to be in fair agreement. Finally, the expected behavior under forced-switching control gives confidence that the experimental testbed is functioning properly.

References

[1] M. Sheppy, L. Gentile-Polese, and S. Gould, “Analysis of Plug Load Capacities and Power Requirements in Commercial Buildings,” report, National Renewable Energy Lab. (Sept. 2014), https://www.nrel.gov/docs/fy14osti/60266.pdf

[2] https://support.egauge.net/support/solutions/articles/25000002525-extending-the-length-of-the-ct-leads

[3] https://srdata.nist.gov/its90/download/type_k.tab

[4] https://www.littelfuse.com/products/temperature-sensors/leaded-thermistors/interchangeable-thermistors/standard-precision-ps.aspx

[5] H. Ebrahimi-Darkhaneh, “Measurement error caused by self-heating in NTC and PTC thermistors”, Texas Instruments Analog Design Journal, Q3, 2019

[6] Richard Bravo, Robert Yinger, Dave Chassin, Henry Huang, Ning Lu, Ian Hiskens, Giri Venkataramanan, FINAL PROJECT REPORT LOAD MODELING TRANSMISSION RESEARCH, CIEE, 2010.

Appendix A: Selected Load

Window Air Conditioner

LW5016

LG
Life's Good

LESS REALLY CAN DO MORE

Say you need just a little extra cooling, but you want a window unit that really works. Maybe one with 5,000 BTUs, dual cooling speeds and a 2-way air direction function that cools up to 150 square feet. Well, we've got it right here. Not a ton of bells and whistles. Just a really great air conditioner.

PERFORMANCE

- 5,000 BTU
- CEER 11.0 / EER 11.2
- Dehumidification (Pts/Hr) 1.3
- Est. Cooling Area 150 sq. ft. (Room size 10' x 15')

FEATURES

- 2 Cooling & Fan Speeds
- 2-Way Air Direction
- Easy-to-Use Installation Kit

2 Cooling & Fan Speeds

Control the temperature of the room at your own pace with our two cooling and fan speeds.

2-Way Air Direction

Enjoy the cool air even more with our two-way air direction. This function allows you to move air left or right, so no hot spot is left untouched.

Easy-to-Use Installation Kit

Get cool with the greatest of ease. This window unit comes with our easy-to-use installation kit so you can set up your air conditioner without straining your back.

PERFORMANCE	
BTU Performance	5,000
CEER	11.0
EER	11.2
ENERGY STAR®	No
Dehumid. (Pts/Hr)	1.3
Dry Air Flow (CFM)	134
dBA Level (Indoor/Outdoor)	52/56
Est. Cooling Area (SQ. FT.)	150
Refrigerant	R410A

FEATURES	
Thermostat Control	Adjustable Thermostat
Air Deflection	2-Way
Fan Speed Cooling	2
Fan Only Speed	2
Compressor	Rotary
In Door Fan Type	Turbo
Type Air Discharge	Top Discharge
Chassis Type	Top Down

MATERIALS/FINISHES	
Available Colors	White

ELECTRICAL RATINGS	
Voltage/60Hz	115
Watts	440
Rated Amps	4.1

DIMENSIONS	
Product (WxHxD)	17 5/16" x 11 1/8" x 14 3/8"
Shipping (WxHxD)	19 5/8" x 13 11/16" x 16 3/4"
Net Weight	40 lbs.
Shipping Weight	44 lbs.

LIMITED WARRANTY	
	1 Year Parts and Labor

UPC CODES	
LW5016	040231 370907

1YR PARTS & LABOR LIMITED IN-HOME SERVICE WARRANTY

115 Volt 15 Amp

LG Electronics U.S.A., Inc.
1000 Sylvan Avenue Englewood Cliffs, NJ 07632
Customer Service and Technical Support: (800) 243-0000
LG.com

Appendix B: Bill of Materials, Mechanical

Mechanical items:

Part Number	Description	Qty
184209643108DS	Pallet racks, 108" beams, 96" struts, wire shelves	5
RSIN30	30" row separators	4
H-1125	Uline 48"x48" pallets	5
LW5016	LG LW5016 air conditioner, 5000 BTU/hr	5
007-F5	Taco water circulation pump	5
110-251F	Taco 3/4" flanges	5
XT-8	Zurn 2 gallon water expansion tank	5
XE20P06PU20U0	Rheem 120V, 1500W, 20 gal. water heater	2
XE30P06PU20U1	Rheem 120V, 1500W, 30 gal. water heater	3
	Dow ScoreBoard blue extruded foam R10 2"x48"x96"	20
625-AF10	Duct fan, 10" dia.	5
FV-4M1- 3/4	Watts, 3/4" Automatic air vent valve	5
	PEX tubing	100'
	2"x4"x8' lumber	var.
	Plywood, 4'x8'	3
	Ball valves	20
HTL 12x12	Water-to-air heat exchangers, Alfa Heating	5
20166-0030	Pressure relief valve, 3/4"	5

Appendix B, continued: Bill of Materials, Electrical

Electrical Items:

Part Number	Description	Qty
	Wire, black, AWG 18 stranded	500'
	Wire, white, AWG 18 stranded	500'
USP 10981	Thermistor, 1/8" NPT sheathed, 10k, Littelfuse	5
PS103J2	Thermistor, interchangeable, 10k, Littelfuse	20
WSU240-0500-R	Triad, 24 Vdc power supplies	6
G4A-1A-E DC24	OMRON, power relay, 24 Vdc	5
721AFMS	Switchcraft panel mount barrel connectors	12
19C7295	DIN rail, 35mm, 1 m length	3
DA10-24F0-0000	Watlow, DIN-a-Mite power controller	5
1182P60	Hammond, power transformer, 2:1, 625 VA	5
XTR110KP	XTR110 voltage-to-current, DIP16 pkg, Texas Instr.	10
480172	Extech AC line splitter	5
EG4130	eGauge Pro power meter, 30 inputs	1
INT16x14x7	Enclosure for eGauge	1
Power Whip	Power connector with plug, 3p + neutral, 120V	1
ACT-20mm-100A	CT sensors, 100A, high accuracy, split core	5
	Ethernet cable	1
6NX6	Power strips, 6 outlet, Tripp-Lite	6
	Computer (Windows 10)	1
	Monitor+kbd+mouse	1
1444-33	Enclosures for TCL local electronics, Hammond	5
PS0SSSSXB	Power entry module, 10A, Tyco Electronics	5
	Terminal blocks	50

Appendix C: Bill of Materials, Data Acquisition System

Data acquisition system:

Part Number	Description	Qty
781161-01	PXIe-1073, 5 peripheral slots, PCIe-8361 adapter, 3m cable	1
763000-01	120V, 10A power cord	1
783800-01	PXIe-6738, 16-bit, 32 Analog Outputs	1
782536-01	SCB-68A shielded connector block	1
157599-01	SCH68-68-A2 shielded cable, 1m	1
783865-01	PXIe-4302, 24-bit, 32 Voltage Inputs	2
783869-01	TB-4302, front-mount terminal blocks	2